UNCONSTRAINED OPTIMIZATION

PAUL SCHRIMPF
OCTOBER 24, 2013
UNIVERSITY OF BRITISH COLUMBIA
ECONOMICS 526

Today's lecture is about unconstrained optimization. If you're following along in the syllabus, you'll notice that we've skipped the fourth topic, eigenvalues and definite matrices. We will cover these things as part of our study of optimization.

1. NOTATION AND DEFINITIONS

An optimization problem refers to finding the maximum or minimum of a function, perhaps subject to some constraints. In economics, the most common optimization problems are utility maximization and profit maximization. Because of this, we will state most of our definitions and results for maximization problems. Of course, we could just as well state each definition and result for a minimization problem by reversing the sign of most inequalities.

In this lecture we will be interested in unconstrained optimization problems such as

$$\max_{x \in U} F(x)$$

where $x \in U \subseteq \mathbb{R}^n$ and $F : U \to \mathbb{R}$. If $F^* = \max_{x \in U} F(x)$, we mean that $F(x) \leq F^*$ for all $x \in U$ and $F(x^*) = F^*$ for some $x^* \in U$.

Definition 1.1. $F^* = \max_{x \in U} F(x)$ is the **maximum** of F on U if $F(x) \leq F^*$ for all $x \in U$ and $F(x^*) = F^*$ for some $x^* \in U$

There may be more than one such x^*. We denote the set of all x^* such that $F(x^*) = F^*$ by $\arg\max_{x \in U} F(x)$ and might write $x^* \in \arg\max_{x \in U} F(x)$, or, if we know there is only one such, x^*, we sometimes write $x^* = \arg\max_{x \in U} F(x)$.

Definition 1.2. $x^* \in U$ is a **maximizer** of F on U if $F(x*) = \max_{x \in U} F(x)$. The set of all maximizers is denoted $\arg\max_{x \in U} F(x)$.

Definition 1.3. $x^* \in U$ is a **strict maximizer** of F on U if $F(x*) > F(x)$ for all $x \in U$ with $x \neq x*$.

Recall the definition of a local maximum from lecture 8.

Definition 1.4. F has a **local maximum** at x if $\exists \delta > 0$ such that $F(y) \leq F(x)$ for all $y \in N_\delta(x) \cap U$. Each such x is called a **local maximizer** of F. If $F(y) < F(x)$ for all $y \neq x$, $y \in N_\delta(x) \cap U$, then we say F has a **strict local maximum** at x.

When we want to be explicit about the distinction between local maximum and the maximum in definition 1.1, we refer to the later as the global maximum.

Example 1.1. Here are some examples of functions from $\mathbb{R}\to\mathbb{R}$ and their maxima and minima.

(1) $F(x) = x^2$ is minimized at $x = 0$ with minimum 0.
(2) $F(x) = c$ has minimum and maximum c. Any x is a maximizer.
(3) $F(x) = \cos(x)$ has maximum 1 and minimum -1. $2\pi n$ for $n \in \mathbb{Z}$ is a maximizer.
(4) $F(x) = \cos(x) + x/2$ has no global maximizer or minimizer, but has many local ones.

2. First order conditions

In lecture 8, we proved that if F has a local maximum at x, then $DF_x = 0$. We restate that theorem and proof here.

Theorem 2.1. *Let $U \subseteq \mathbb{R}^n$, $F : U\to\mathbb{R}$, and suppose F has a local maximum or minimum at x, F is differentiable at x, and $x \in$ interior(U) . Then $DF_x = 0$.*

Proof. We will write the proof for when F has a local maximum at x. The exact same reasoning works when F has a local minimum.

Since x is in the interior of U, we can choose $\delta > 0$ such that $N_\delta(x) \subset U$. Since x is a local maximum we can also choose x such that $F(x) \geq F(y)$ for all $y \in N_\delta(x)$. Since F is differentiable, we can write

$$\frac{F(x+h) - F(x)}{\|h\|} = \frac{DF_x h + r(x,h)}{\|h\|}$$

where $\lim_{h\to 0} \frac{|r(x,h)|}{\|h\|} = 0$. Let $h = tv$ for some $v \in \mathbb{R}^n$ with $\|v\| = 1$ and $t \in \mathbb{R}$. If $DF_x v > 0$, then for $t > 0$ small enough, we would have $\frac{F(x+tv)-F(x)}{|t|} = DF_x v + \frac{r(x,tv)}{|t|} > DF_x v/2 > 0$ and $F(x+tv) > F(x)$ in contradiction to x being a local maximum. Similarly, if $DF_v v < 0$ then for $t < 0$ and small, we would have $\frac{F(x+tv)-F(x)}{|t|} = -DF_x v + \frac{r(x,tv)}{|t|} > -Df_x v/2 > 0$ and $F(x+tv) > F(x)$. Thus, it must be that $DF_x v = 0$ for all v, i.e. $DF_x = 0$. □

The first order condition is the fact that $DF_x = 0$ is a necessary condition for x to be a local maximizer or minimizer of F.

Definition 2.1. Any point such that $DF_x = 0$ is call a **critical point** of F.

If F is differentiable, F cannot have local minima or maxima (=local extrema) at non-critical points. F might have a local extrema its critical points, but it does not have to. Consider $F(x) = x^3$ $F'(0) = 0$, but 0 is not a local maximizer or minimizer of F. Similarly, if $F : \mathbb{R}^2\to\mathbb{R}$, $F(x) = x_1^2 - x_2^2$, then $DF_0 = 0$, but 0 is not a local minimum or maximum of F.

3. Second order conditions

To determine whether a given critical point is a local minimum or maximum or neither we can look at the second derivative of the function. Let $F : \mathbb{R}^n\to\mathbb{R}$ and suppose x^* is a critical point. Then $DF_{x^*} = 0$. To see if x^* is a local maximum, we need to look at

$F(x^* + h)$ for small h. If F is twice continuously differentiable, we can take a second order Taylor expansion of F around x^*.

$$F(x^* + h) = F(x^*) + DF_{x^*}h + \frac{1}{2}h^T D^2 F_{x*} h + r(x^*, h)$$

where $D^2F_{x^*}$ is the matrix of F's second derivatives. It is called the **Hessian** of F.

$$D^2F_{x^*} = \begin{pmatrix} \frac{\partial^2 F}{\partial x_1^2} & \cdots & \frac{\partial^2 F}{\partial x_1 \partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 F}{\partial x_1 \partial x_n} & \cdots & \frac{\partial^2 F}{\partial x_n}^2 \end{pmatrix}.$$

Since x^* is a critical, point $DF_{x^*} = 0$, so

$$F(x^* + h) - F(x^*) = \frac{1}{2}h^T D^2 F_{x*} h + r(x^*, h).$$

We can see that x^* is a local maximum if

$$\frac{1}{2}h^T D^2 F_{x*} h + r(x^*, h) \leq 0$$

for all $h \neq 0$, $\|h\| < \delta$ for some $\delta > 0$. We know that $r(x^*, h)$ is small so, we expect that the above inequality will be true if $h^T D^2 F_{x*} h \leq 0$ for all $h \neq 0$. The Hessian, D^2F_{x*} is just some symmetric n by n matrix, and $h^T D^2 F_{x*} h$ is a quadratic form in h. This motivates the following definition.

Definition 3.1. Let A be a symmetric matrix, then A is

- **Negative definite** if $x^T Ax < 0$ for all $x \neq 0$
- **Negative semi-definite** if $x^T Ax \leq 0$ for all $x \neq 0$
- **Positive definite** if $x^T Ax > 0$ for all $x \neq 0$
- **Positive semi-definite** if $x^T Ax \geq 0$ for all $x \neq 0$
- **Indefinite** if $\exists\ x_1$ s.t. $x_1^T Ax_1 > 0$ and some other x_2 such that $x_2^T Ax_2 < 0$.

In the next section we will derive some conditions on A that ensure it is negative (semi-)definite. For now, just observe that if D^2F_{x*} is negative semi-definite, then x^* must be a local maximum. If $D^2F_{x^*}$ is negative definite, then x^* is a strict local maximum. The following theorem restates the results of this discussion.

Theorem 3.1. *Let $F : U \to \mathbb{R}$ be twice continuously differentiable on U and let x^* be a critical point in the interior of U. If*

(1) *The Hessian, $D^2F_{x^*}$ is negative definite, then x^* is a strict local maximizer.*
(2) *The Hessian, $D^2F_{x^*}$ is positive definite, then x^* is a strict local minimizer.*
(3) *The Hessian, $D^2F_{x^*}$ is indefinite, x^* is neither a local min nor a local max.*
(4) *The Hessian is positive or negative semi-definite, then x^* could be a local maximum, minimum, or neither.*

Proof. We will only prove the first case. The second and third cases can be proven similarly. We will go over some examples of the fourth case.

The main idea of the proof is contained in the discussion at the start of this section. The only tricky part is carefully showing that $r(x,h)$ is small enough to ignore. As in the discussion preceding the theorem, x^* is a local minimizer if

$$F(x^*+h) - F(x^*) = \frac{1}{2}h^T D^2 F_{x^*} h + r(x^*,h) \leq 0.$$

We can rewrite this as

$$h^T D^2 F_{x^*} h + r(x^*,h) = h^T D^2 F_{x^*} h + h^T h \frac{r(x^*,h)}{\|h\|^2}$$

Factoring, we have

$$h^T D^2 F_{x^*} h + h^T h \frac{r(x^*,h)}{\|h\|^2} = h^T \left(D^2 F_{x^*} + \frac{r(x^*,h)}{\|h\|^2} \right) h.$$

From our theorem on Taylor series, we know $\lim_{h\to 0} \frac{r(x^*,h)}{\|h\|^2} = 0$. Then for any $\epsilon > 0$ $\exists$ $\delta > 0$ such that if $\|h\| < \delta$, then $\frac{|r(x^*,h)|}{\|h\|^2} < \epsilon$. So,

$$\begin{aligned} h^T \left(D^2 F_{x^*} + \frac{r(x^*,h)}{\|h\|^2} \right) h &\leq h^T \left(D^2 F_{x^*} + \epsilon \right) h \\ &\leq h^T D^2 F_{x*} h + \epsilon \|h\|^2 \end{aligned}$$

Let $t = \|h\|$ and $q = \frac{h}{\|h\|}$, so $h = tq$. Then we have

$$h^T \left(D^2 F_{x^*} + \frac{r(x^*,h)}{\|h\|^2} \right) h \leq t^2 (q^T D^2 F_{x*} q) + \epsilon t^2$$

We can pick $\epsilon < \inf_{\|q\|=1} |q^T D^2 F_{x*} q|$. The set $\{q : \|q\| = 1\}$ is compact so there is some q that achieves this minimum, and $q^T D^2 F_{x^*} q < 0$, so $\epsilon > 0$.[1] Then for all $|t| < \delta$ and $\|q\| = 1$,

$$t^2 (q^T D^2 F_{x*} q) + \epsilon t^2 < t^2 (q^T D^2 F_{x*} q) + |(q^T D^2 F_{x*} q)| t^2 < 0$$

where the last inequality follows from $q^T D^2 F_{x^*} q < 0$. Thus,

$$F(x^*+h) - F(x^*) < 0$$

for all $\|h\| < \delta$. □

When the Hessian is not positive definite, negative definite, or indefinite, the result of this theorem is ambiguous. Let's go over some examples of this case.

Example 3.1. $F : \mathbb{R} \to \mathbb{R}$, $F(x) = x^4$. The first order condition is $4x^3 = 0$, so $x^* = 0$ is the only critical point. The Hessian is $F''(x) = 12x^2 = 0$ at x^*. However, x^4 has a strict local minimum at 0.

[1] If we didn't have compactness, we might have ended up with $\epsilon = 0$, which is not allowed.

Example 3.2. $F : \mathbb{R}^2 \to \mathbb{R}$, $F(x_1, x_2) = -x_1^2$. The first order condition is $DF_x = (-2x_1, 0) = 0$, so the $x_1^* = 0$, $x_2^* \in \mathbb{R}$ are all critical points. The Hessian is

$$D^2F_x = \begin{pmatrix} -2 & 0 \\ 0 & 0 \end{pmatrix}$$

This is negative semi-definite because $h^T D^2F_x h = -2h_1^2 \leq 0$. Also, graphing the function would make it clear that $x_1^* = 0$, $x_2^* \in \mathbb{R}$ are all (non-strict) local maxima.

Example 3.3. $F : \mathbb{R}^2 \to \mathbb{R}$, $F(x_1, x_2) = -x_1^2 + x_2^4$. The first order condition is $DF_x = (-2x_1, 4x_2^3) = 0$, so the $x^* = (0, 0)$ is a critical point. The Hessian is

$$D^2F_x = \begin{pmatrix} -2 & 0 \\ 0 & 12x_2^2 \end{pmatrix}$$

This is negative semi-definite at 0 because $h^T D^2F_0 h = -2h_1^2 \leq 0$. However, 0 is not a local maximum because $F(0, x_2) > F(0, 0)$ for any $x_2 \neq 0$. 0 is also not a local minimum because $F(x_1, 0) < F(0, 0)$ for all $x_1 \neq 0$.

In each of these examples, the second order condition is inconclusive because $h^T D^2F_{x^*} h = 0$ for some h. In these cases we could determine whether x^* is a local maximum, local minimum, or neither by either looking at higher derivatives of F at x^*, or look at D^2F_x for all x in a neighborhood of x^*. We will not often encounter cases where the second order condition is inconclusive, so we will not study these possibilities in detail.

The converse of theorem 3.1 is nearly true. If x^* is a local maximizer, then $D^2F_{x^*}$ must be negative semi-definite.

Theorem 3.2. *Let $F : U \to \mathbb{R}$ be twice continuously differentiable on U and let x^* in the interior of U be a local maximizer (or minimizer) of F. Then $DF_{x^*} = 0$ and $D^2F_{x^*}$ is negative (or positive) semi-definite.*

Proof. Using the same notation and setup as in the proof of theorem 3.1,

$$\begin{aligned} 0 > F(x^* + tq) - F(x^*) =& t^2 \frac{1}{2} q^T D^2F_{x^*} q + r(x, tq) \\ 0 >& t^2 \left(\frac{1}{2} q^T D^2F_{x^*} q + \frac{r(x, tq)}{t^2} \right) \end{aligned}$$

Because $\lim_{t \to 0} \frac{r(x,tq)}{t^2} = 0$, for any $\epsilon > 0$ $\exists \delta > 0$ such that if $|t| < \delta$, then

$$\begin{aligned} 0 >& t^2 \left(\frac{1}{2} q^T D^2F_{x^*} q + \frac{r(x, tq)}{t^2} \right) \geq t^2 \left(\frac{1}{2} q^T D^2F_{x^*} q - \epsilon \right) \\ t^2 \epsilon >& \frac{1}{2} t^2 q^T D^2F_{x^*} q \end{aligned}$$

This is only possible for all $\epsilon > 0$ if $q^T D^2F_{x^*} q \leq 0$. □

I do not expect you to remember the proofs of the last two theorems. However, it is important to remember the second order condition, and know how to check it. To check the second order condition, we need some practical way to tell whether a matrix is positive or negative (semi-)definite.

4. DEFINITE MATRICES

The second order condition says that a critical point is a local maximum if the Hessian is negative definite. In this section we will develop some conditions for whether a matrix is negative definite. Let's start with the simplest case where A is a one by one matrix. Then $x^TAx = x^2a$ is negative definite if and only if $a < 0$. A is negative semi-definite if $a \leq 0$. What if A is a two by two symmetric matrix? Then,

$$\begin{aligned} x^TAx =& \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}^T \begin{pmatrix} a & b \\ b & c \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \\ =& ax_1^2 + 2bx_1x_2 + cx_2^2 \end{aligned}$$

Completing the square we get

$$\begin{aligned} x^TAx =& a\left(x_1 + \frac{2b}{a}x_1x_2 + \frac{b^2}{a^2}x_2^2\right) - \frac{b^2}{a}x_2^2 + cx_2^2 \\ =& a\left(x_1 + \frac{b}{a}x_2\right)^2 + \frac{ac - b^2}{a}x_2^2 \end{aligned}$$

Thus $x^TAx < 0$ for all $x \neq 0$ if $a < 0$ and $\frac{ac-b^2}{a} < 0$, i.e. $ac - b^2 > 0$. Notice that $ac - b^2 = \det A$. So A is negative definite if $a < 0$ and $\det A > 0$. If we wanted A to be positive definite, we would need $a > 0$ and $\det A > 0$. For semi-definite we get the same thing with weak instead of strict inequalities. It turns out that in general this sort of pattern continues. We can determine whether A is negative definite by looking at the determinants of certain submatrices of A.

Definition 4.1. Let A by an n by n matrix. The k by k submatrix

$$A_k = \begin{pmatrix} a_{11} & \cdots & a_{1k} \\ \vdots & & \vdots \\ a_{k1} & \cdots & a_{kk} \end{pmatrix}$$

is the k**th leading principal submatrix** of A. The determinant of A_k is the kth order **leading principal minor** of A.

Theorem 4.1. *Let A be an n by n symmetric matrix. Then*

1. *A is positive definite if and only if all n of its leading principal minors are strictly positive.*
2. *A is positive semi-definite if and only if all n of its leading principal minors are weakly positive.*
3. *A is negative definite if and only if all n of its leading principal minors alternate in sign as follows:* $\det A_1 < 0$, $\det A_2 > 0$, $\det A_3 < 0$, *etc.*
4. *A is negative semi-definite if and only if all n of its leading principal minors weakly alternate in sign as follows:* $\det A_1 \leq 0$, $\det A_2 \geq 0$, $\det A_3 \leq 0$, *etc*
5. *A is indefinite if and only if none of the five above cases hold, and* $\det A_k \neq 0$ *for at least one k.*

The proof of this theorem is a bit tedious, so we will skip it. There is a proof in chapter 16.4 of Simon and Blume. Figure 1 shows each of the five types of definiteness.

FIGURE 1. Definite quadratic forms

Positive definite $x^2 + y^2$

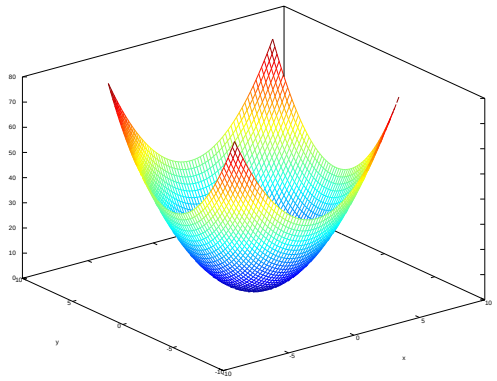

Negative definite $-x^2 - y^2$

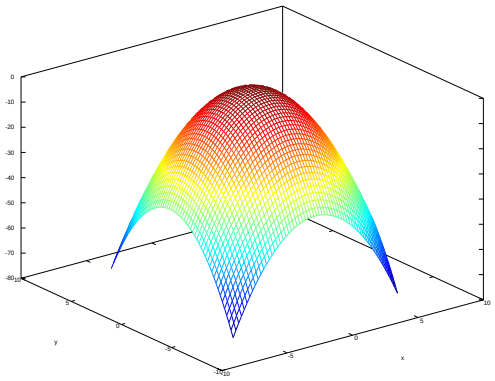

Positive semi-definite x^2

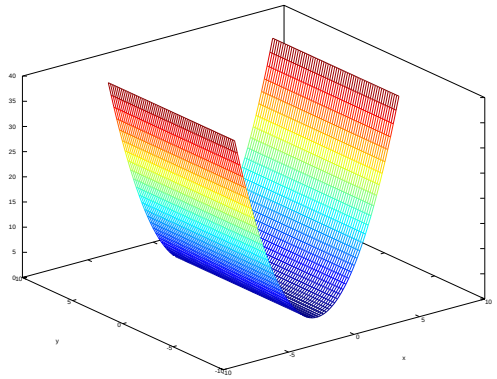

Negative semi-definite $-y^2$

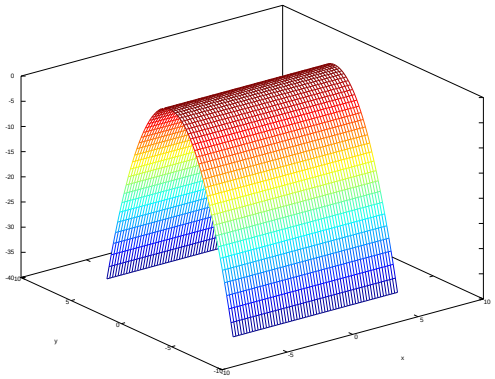

Indefinite $x^2 - y^2$

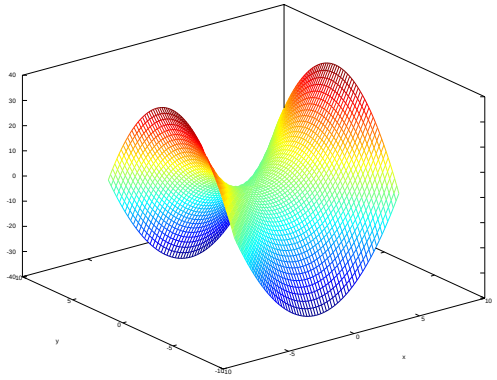

4.1. **Eigenvectors and eigenvalues.** Another way to check whether a matrix is negative definite is by looking at the matrix's eigenvalues. Eigenvalues are of interest in their own right as well because eigenvalues have many other uses, such as determining the stability of systems of difference and differential equations.

Definition 4.2. If A is an n by n matrix, λ is a scalar, $v \in \mathbb{R}^n$ with $\|v\| = 1$, and

$$Av = \lambda v$$

then λ is a **eigenvalue** of A and v is an **eigenvector**.

If v is an eigenvector of A with eigenvalue λ, then

$$A(tv) = \lambda(tv)$$

for any $t \neq 0$. Thus, the requirement that $\|v\| = 1$ is just a normalization to pin down v. There are a few equivalent ways of defining eigenvalues, some of which are given by the following lemma.

Lemma 4.1. *Let A be an n by n matrix and λ a scalar. Each of the following are equivalent.*

(1) *λ is a eigenvalue of A.*
(2) *$A - \lambda I$ is singular.*
(3) *$\det(A - \lambda I) = 0$.*

Proof. We know that (2) and (3) from our results on systems of linear equations and matrices. Also, if $A - \lambda I$ is singular, then the null space of $A - \lambda I$ contains non-zero vectors. Choose $v \in \mathcal{N}(A - \lambda I)$ such that $v \neq 0$. Then $(A - \lambda I)v = 0$, so $A(v/\|v\|) = \lambda(v/\|v\|)$ as in the definition of eigenvalues. □

The function $\chi_A : \mathbb{R} \to \mathbb{R}$ defined by $\chi_A(x) = \det(A - xI)$ is called the **characteristic polynomial** of A. It will be a polynomial of order n. You may know from some other math course that polynomials of degree n have n roots, some of which might be not be real and some of which might not be distinct.

Suppose that A has $m \leq n$ distinct real eigenvalues, $\lambda_1, ..., \lambda_m$. Associated with each eigenvalue is a linear subspace of eigenvalues, $\{v \in \mathbb{R}^n : Av = \lambda_i v\} = \mathcal{N}(A - \lambda_i I)$. By definition of eigenvalues each of these subspaces has dimension at least 1. Some of them may be of larger dimension, but they can be of dimension at most n. For each $\mathcal{N}(A - \lambda_i I)$ there $\exists$ an orthonormal basis of eigenvectors, i.e. $v_i^1, ..., v_i^{k_i}$ such that $Av_i^j = \lambda_i v_i^j$, $\left\|v_i^j\right\| = 1$ and $(v_i^j)^T v_i^l = 0$ for $j \neq l$.

Lemma 4.2. *Let A be an n by n matrix with m distinct real eigenvalues, and let $v_i^1, ..., v_i^k$ be an orthonormal basis for $\mathcal{N}(A - \lambda_i I)$. Then $\{v_1^1, ..., v_1^{k_1}, ..., v_m^1, ..., v_m^{k_m}\}$ are linearly independent.*

Proof. Suppose the eigenvectors are not linearly independent. Then we could write[2]

$$v_1^1 = c_1^2 v_1^2 + ... + c_m^k v_m^k.$$

[2]Possibly we cannot do this for v_1, but we can for some v_i. Without loss of generality, assume that we can for v_1.

with at least one $c_i^j \neq 0$. By the definition of eigenvalues and eigenvectors, we have

$$\begin{aligned} Av_1 =& \lambda_1 v_1 \\ A(c_1^2 v_1^2 + ... + c_m^k v_m^k) =& \lambda_1 \left(c_1^2 v_1^2 + ... + c_m^k v_m^k) \right) \\ c_1^2 \lambda_1 v_1^2 + ... + c_m^k \lambda_m v_m^k =& \lambda_1 \left(c_1^2 v_2 + ... + c_m^k v_m^k \right) \\ c_2^1(\lambda_2 - \lambda_1) v_2^1 + ... + c_m^k (\lambda_m - \lambda_1) v_m =& 0 \end{aligned}$$

If $c_2^1 = \cdots = c_m^k$, then the original equation becomes

$$v_1^1 = c_1^2 v_1^2 + \cdots + c_1^{k_1} v_1^{k_1}$$

with one of these c's non-zero. That would contradict the way $v_1^1, ...v_1^{k_1}$ were constructed to be linearly independent. Hence, $c_i^j \neq 0$ for at least one $i > 1$. Since the λ_i are distinct, $\lambda_k - \lambda_1 \neq 0$ for any k, so this says that $v_2^1, ..., v_m^{k_m}$ are linearly dependent as well. We can repeat this argument to show that $v_3^1, ..., v_m^{k_m}$ are linearly dependent, and then repeat it again and again and eventually show that v_{m-1} and v_m are linearly dependent. Then $v_{m-1} = bv_m$ for some $b \neq 0$, and

$$Av_{m-1} = \lambda_{m-1} v_{m-1} = \lambda_{m-1}(bv_m)$$

and

$$Av_{m-1} = Abv_m = b\lambda_m v_m.$$

This implies that $\lambda_m = \lambda_{m-1}$, contrary to the eigenvalues being distinct. Therefore, the collection of eigenvectors must be linearly independent. □

Since the eigenvectors are linearly independent, in $\mathbb{R}^n$ there can be at most n of them. However, there are some matrices that have strictly fewer eigenvectors.

To say much more about eigenvalues we must impose the extra condition that A is symmetric. Fortunately, Hessian matrices are symmetric, so for the current purposes, imposing symmetry is not a problem. Since A is symmetric, $A = A^T$. Also, by the definition of transpose, $\langle Ax, y \rangle = \langle x, A^T y \rangle$, or specializing to $\mathbb{R}^n$, $(Ax)^T y = xA^T y = xAy$, where the last equality comes from symmetry of A. By the fundamental theorem of algebra (which says that polynomials have (possibly complex) roots), $\exists$ at least one λ_1 and v_1 such that

$$Av_1 = \lambda_1 v_1.$$

A short argument using complex numbers (which we have not covered) and the symmetry of A shows that λ_1 must be real. If we then consider A as a symmetric linear operator on the $n-1$ dimensional space $\text{span}\{v_1\}^\perp$, we can repeat this argument to show that there is a second (not necessarily distinct) eigenvalue λ_2 and associated eigenvector v_2 with $v_2 \perp v_1$. If we repeat this n times, we will construct n eigenvalues (some of which may be the same) and n orthogonal eigenvectors. Moreover, we can rescale the eigenvectors so that $\|v_i\| = 1$ for each i. These n eigenvectors form an orthonormal basis for $\mathbb{R}^n$. You can imagine them as some rotation of the usual axes.

Next, we can write A in terms of its eigenvectors and eigenvalues. If we make Λ a diagonal matrix consisting of the eigenvalues of A and V a matrix whose columns are the eigenvectors of A, then

$$AV = V\Lambda.$$

Moreover, V is an orthogonal matrix, so $V^{-1} = V^T$. This relationship is called the eigendecomposition of A.

Theorem 4.2 (Eigendecomposition)**.** *Let A be an n by n symmetric matrix, then A has n (not necessarily distinct) eigenvalues and*

$$A = V\Lambda V^T$$

where Λ is the diagonal matrix consisting of the eigenvalues of A and the columns of V are the associated eigenvectors of A, and V is an orthonormal matrix.

Comment 4.1. There are non-symmetric matrices that can not be decomposed into eigenvalues and eigenvectors, for example $\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$.

There are other non-symmetric matrices that can be eigendecomposed, for example $\begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 1 & 0 & 1 \end{pmatrix}$ can be eigendecomposed. Any square matrix with $A^TA = AA^T$ can be eigendecomposed.

Using the eigendecomposition, we can relate eigenvalues to the definiteness of a matrix.

Theorem 4.3. *If A is an n by n symmetric non-singular matrix with eigenvalues $\lambda_1, ..., \lambda_n$, then*

(1) *$\lambda_i > 0$ for all i, iff A is positive definite,*
(2) *$\lambda_i \geq 0$ for all i, iff A is positive semi-definite,*
(3) *$\lambda_i < 0$ for all i, iff A is negative definite,*
(4) *$\lambda_i \leq 0$ for all i, iff A is negative semi-definite,*
(5) *if some $\lambda_i > 0$ and some $\lambda_j < 0$, then A is indefinite.*

Proof. Let $A = V\Lambda V^T$ be the eigendecomposition of A. Let $x \neq 0 \in \mathbb{R}^n$. Then $x^TV = z^T \neq 0$ because V is nonsingular. Also,

$$x^TAx = x^TV\Lambda V^Tx = z^T\Lambda z.$$

Since Λ is diagonal we have

$$z^T\Lambda z = z_1^2\lambda_1 + ... + z_n^2\lambda_n$$

Thus, $x^TAx = z^T\Lambda z > 0$ for all x iff $\lambda_i > 0$ for each i. The other parts of the theorem follow similarly. □

4.2. **Other facts about eigendecomposition.** The results in this subsection are not essential for the rest of this lecture, but they are occasionally useful.

If a matrix is singular, then $\exists\ x \neq 0$ such that $Ax = 0$. By definition, then $x/\|x\|$ is an eigenvector and 0 is an eigenvalue of A. Conversely, if 0 is an eigenvalue of A, then $\exists x$ with $\|x\| = 1$ such that $Ax = 0$. Then A is singular.

Lemma 4.3 (Eigenvalues and invertibility)**.** *A is singular iff 0 is an eigenvalue of A.*

Recall that we defined the norm of a linear transformation as

$$\|A\|_{BL} = \sup_{x:\|x\|=1} \|Ax\| .$$

When A is symmetric (so that its eigenvalues exist), this norm is equal to the largest eigenvalue of A.

Lemma 4.4 (Eigenvalues and operator norm)**.** *Suppose the $n \times n$ matrix A has n eigenvalues, $\lambda_1, ..., \lambda_n$, then $\|A\|_{BL} = \max_{1 \leq i \leq n} \lambda_i$.*

We saw that definiteness is related to both determinant and eigenvalues. It should come as no suprise that determinants and eigenvalues are directly related. The determinant of a matrix is equal to the product of its eigenvalues.

Lemma 4.5 (Eigenvalues and determinants)**.** *If A has eigenvalues $\lambda_1, ..., \lambda_n$, such that*

$$\det(A - tI) = \prod_{i=1}^{n} (t - \lambda_i)$$

then $\det(A) = \prod_{i=1}^{n} \lambda_i$.

5. Global maximum and minimum

The second order condition (3.1) along with the first order condition gives a nice way of finding local maxima of a function, but what about the global maximum? In general, you could use the first and second order conditions to find each of the local maxima of F in the interior of U. You then must compare each of these local maxima with each other and the value of F on the boundary of U to find the global maximum of F on U. If there are many local maxima or the boundary of U is complicated, this procedure can be quite complicated. It would be nice if there was some simpler necessary and sufficient for a global maximum. Unfortunately, there is no such general necessary and sufficient condition. There is, however, a sufficient condition that is sometimes useful.

Definition 5.1. Let $f : U \to \mathbb{R}$. f is **convex** if for all $x, y \in U$ with $\ell(x, y) \subseteq U$ we have $f(tx + (1-t)y) \leq tf(x) + (1-t)f(y)$ for all $t \in [0, 1]$.

Equivalently, F is convex if the set $\{(y, x) : x \in U, y \geq f(x)\} \subseteq U \times \mathbb{R}$ is convex. This set is called the epigraph of f. If you draw the graph of the function, the epigraph is the set of points above the function.

Definition 5.2. Let $f : U \to \mathbb{R}$. f is **concave** if for all $x, y \in U$ with $\ell(x, y) \subseteq U$ we have $f(tx + (1-t)y) \geq tf(x) + (1-t)f(y)$ for all $t \in [0, 1]$.

Equivalently, F is concave if the set $\{(y, x) : x \in U, y \leq f(x)\}$ is convex.

Theorem 5.1. *Let $F : U \to \mathbb{R}$ be twice continuously differentiable and $U \subseteq \mathbb{R}^n$ be convex. Then,*

(1) *For maximization:*
   (a) *The following three conditions are equivalent:*

(i) *F is a concave function on U*
(ii) *$F(y) - F(x) \leq DF_x(y - x)$ for all $x, y \in U$,*
(iii) *D^2F_x is negative semi-definite for all $x, y \in U$*

(b) *If F is a concave function on U and $DF_{x^*} = 0$ for some $x^* \in U$, then x^* is the global maximizer of F on U.*

(2) *For minimization:*

(a) *The following three conditions are equivalent:*
(i) *F is a convex function on U*
(ii) *$F(y) - F(x) \geq DF_x(y - x)$ for all $x, y \in U$,*
(iii) *D^2F_x is positive semi-definite for all $x, y \in U$*

(b) *If F is a convex function on U and $DF_{x^*} = 0$ for some $x^* \in U$, then x^* is the global minimizer of F on U.*

This theorem is often summarized by saying that convex minimization problems (i.e. the set U and function F are both convex) have unique solutions. We will not prove this theorem until after covering convexity in more detail, which we may or may not have time for. An important fact about convex optimization problems is that their solutions can be efficiently computed. A general optimization problem is what is known as NP-hard, which in practice means that you can write down what looks like a reasonable not too large optimization problem whose solution takes prohibitively long to compute. In contrast convex optimization problems can be solved in polynomial time. In particular, if f is convex, x is n dimensional, the steps needed to compute $f(x)$ is M, and you want to find $\hat{x}$ such that $|f(\hat{x}) - f(x^*)| < \epsilon$, then there are algorithms that can find $\hat{x}$ in $O\left(n(n^3 + M)\log(1/\epsilon)\right)$ steps. In fact, there are algorithms for special types of convex problems that can solve them nearly as fast as a least squares problem or a linear program. In practice, this means that most convex problems can be solved quickly with up to about a one thousand variables, and some convex problems can be solved very quickly even with tens or hundreds of thousands of variables. The discovery that convex optimization problems can be solved efficiently was fairly recent. The main theoretical breakthroughs began in the late eighties and early nineties, and it remains an active area of research.

6. APPLICATIONS

6.1. **Profit maximization.**

6.1.1. *Competitive multi-product firm.* Suppose a firm has produces k goods using n inputs with production function $f : \mathbb{R}^n \to \mathbb{R}^k$. The prices of the goods are p, and the prices of the inputs are w, so that the firm's profits are

$$\Pi(x) = p^T f(x) - w^T x.$$

The firm chooses x to maximize profits.

$$\max_x p^T f(x) - w^T x$$

The first order condition is

$$p^T Df_{x^*} - w = 0.$$

or without using matrices,

$$\sum_{j=1}^{k} p_j \frac{\partial f_j}{\partial x_i}(x^*) = w_i$$

for $i = 1, ..., n$. The second order condition is that

$$D[p^T Df]_{x^*} = \begin{pmatrix} \sum_{j=1}^{k} p_j \frac{\partial^2 f_j}{\partial x_1^2}(x^*) & \cdots & \sum_{j=1}^{k} p_j \frac{\partial^2 f_j}{\partial x_1 \partial x_n}(x^*) \\ \vdots & & \vdots \\ \sum_{j=1}^{k} p_j \frac{\partial^2 f_j}{\partial x_1 \partial x_n}(x^*) & \cdots & \sum_{j=1}^{k} p_j \frac{\partial^2 f_j}{\partial x_n^2}(x^*) \end{pmatrix}$$

must be negative semidefinite.

6.1.2. *Multi-product monopolist.* Consider the same setup as before, but now with a monopolist who recognizes that prices depend on output. Let the inverse demand function be $P(q)$ where $P : \mathbb{R}^k \to \mathbb{R}^k$. Now the firm's problem is

$$\max_x P(f(x))^T f(x) - w^T x$$

The first order condition is

$$Df_{x^*}^T DP_{f(x^*)}^T f(x^*) + P(f(x^*))Df_{x^*} - w = 0$$

or without matrices,

$$\sum_{j=1}^{k} \left(\sum_{l=1}^{k} \frac{\partial P_j}{\partial q_l}(f(x^*)) \frac{\partial f_l}{\partial x_i}(x^*) f_l(x^*) + P_j(f(x^*)) \frac{\partial f_j}{\partial x_i}(x^*) \right) = w_i$$

We can get something a bit more interpretable by writing this in terms of elasticities. Recall that the elasticity of demand for good l with respect to the price of good j is $(\epsilon_j^l)^{-1} = \frac{\partial P_j}{\partial q_l} \frac{q_l}{P_l}$. Then,

$$\sum_{j=1}^{k} P_j(f(x^*)) \left(\sum_{l=1}^{k} (\epsilon_j^l)^{-1} \frac{\partial f_l}{\partial x_i}(x^*) + \frac{\partial f_j}{\partial x_i}(x^*) \right) = w_i$$

$$\sum_{j=1}^{k} P_j(f(x^*)) \left[\frac{\partial f_j}{\partial x_i}(x^*) \left(1 + (\epsilon_j^j)^{-1} \right) + \sum_{l \neq j} (\epsilon_j^l)^{-1} \frac{\partial f_l}{\partial x_i}(x^*) \right] = w_i$$

There is a way to compare the price and quantity produced by the monopolist to the competitive firm. There are also things that can be said about the comparison between a single product monopolist ($k = 1$) vs a multi-product monopolist. It might be interesting to derive some of these results. To begin with let $k = 1$ and compare the monopolist to the competitive firm. Under some reasonable assumptions on f and P, you can show that $x^m < x^c$ where x^m is x^* for the monopolist and x^c is x^* for the competitive firm. You can also show that $p^m > p^c$ and that $|\epsilon_1^1| < 1$ for the monopolist. (Although I left it out of the notation above, ϵ depends on x).