MA Functional Analysis 3421
2018 Lecture Notes

John Stalker[*], TCD

Revision: 17.0
Date: 2018-11-14 14:03:57+00
Check `http://www.maths.tcd.ie/~stalker/3421/` for the most recent version.

Contents

1 Inequalities

1.1 Finite Sums

Here and in what follows the real numbers are denoted by $\mathbf{R}$ and the positive reals by $\mathbf{R}^+$. The complex numbers are $\mathbf{C}$ and $\mathbf{K}$ can be taken to be either $\mathbf{R}$ or $\mathbf{C}$. Suppose $n \geq 1$, $x = (\xi_1, \ldots \xi_n) \in \mathbf{R}$, $y = (\eta_1, \ldots, \eta_n) \in \mathbf{R}^+$ and φ is strictly convex. Then

$$\varphi\left(\frac{\sum_{j=1}^{n} \omega_j \xi_j}{\sum_{j=1}^{n} \omega_j}\right) \leq \frac{\sum_{j=1}^{n} \omega_j \varphi\left(\xi_j\right)}{\sum_{j=1}^{n} \omega_j}$$

with equality if and only if all ξ's are equal. This is clear if $n = 1$. For $n > 1$ it is proved by induction.

It's convenient to introduce the quantities

$$\alpha_m = \sum_{j=1}^{m} \omega_j \xi_j, \quad \beta_m = \sum_{j=1}^{m} \omega_j,$$

$$\gamma_m = \sum_{j=1}^{m} \omega_j \varphi\left(\xi_j\right).$$

Our induction hypothesis is then

$$\varphi((\alpha_k/\beta_k) \leq \gamma_k/\beta_k$$

and we wish to prove the same with k replaced by $k+1$. Let

$$s = \frac{\beta_k}{\beta_{k+1}}, \quad t = \frac{\omega_{k+1}}{\beta_{k+1}}, \quad \mu = \frac{\alpha_k}{\beta_k}, \quad \nu = \xi_{k+1}.$$

Then $s, t \geq 0$ and $s + t = 1$ so, by the definition of strict convexity,

$$\varphi(s\mu + t\nu) \leq s\varphi(\mu) + t\varphi(\nu).$$

with equality if and only if $\mu = \nu$. Now

$$s\mu = \frac{\alpha_k}{\beta_{k+1}}, \quad t\nu = \frac{\omega_{k+1}\xi_{k+1}}{\beta_{k+1}}, \quad s\mu + t\nu = \frac{\alpha_{k+1}}{\beta_{k+1}},$$

and hence

$$\varphi\left(\frac{\alpha_{k+1}}{\beta_{k+1}}\right) \leq \frac{\beta_k\varphi\left(\alpha_k/\beta_k\right) + \omega_k\varphi\left(\xi_{k+1}\right)}{\beta_{k+1}}$$

with equality if and only if

$$\frac{\alpha_k}{\beta_k} = \xi_{k+1}.$$

Combining this with the induction hypothesis,

$$\varphi\left(\frac{\alpha_{k+1}}{\beta_{k+1}}\right) \leq \frac{\gamma_k + \omega_k\varphi\left(\xi_{k+1}\right)}{\beta_{k+1}} = \frac{\gamma_{k+1}}{\beta_{k+1}}$$

with equality if and only if both

$$\frac{\alpha_k}{\beta_k} = \xi_{k+1}.$$

and $\xi_1 = \xi_2 = \cdots = \xi_k$. This happens if and only if $\xi_1 = \xi_2 = \cdots = \xi_{k+1}$, so the inductive proof is complete.

We are mostly interested in the special case

$$\sum_{j=1}^{n} \omega_j = 1,$$

in which case the inequality simplifies to

$$\varphi\left(\sum_{j=1}^{n} \omega_j \xi_j\right) \leq \sum_{j=i}^{n} \omega_j \varphi\left(\xi_j\right).$$

This is known as *Jensen's inequality*. Clearly we can allow some of the ω's to be zero, but nothing is really gained by doing so.

Taking $\xi_j = \log\theta_j$ with $\theta_j \in \mathbf{R}^+$ and $\varphi = \exp$ in Jensen's inequality shows that

$$\prod_{j=1}^{n} \theta_j^{\omega_j} \leq \sum_{j=1}^{n} \omega_j \theta_j$$

with equality if and only if all a's are equal. We can allow some or all of the θ's to be zero, but nothing is really gained by doing so. The special case $\omega_j = 1/n$,

$$\left(\prod_{j=1}^{n} \theta_j\right)^{1/n} \leq \frac{\sum_{j=1}^{n} \theta_j}{n},$$

is known as the *arithmetic-geometric mean inequality.*

Applying the inequality

$$\prod_{j=1}^{n} \theta_j^{\omega_j} \leq \sum_{j=1}^{n} \omega_j \theta_j$$

to

$$\theta_j = \frac{\sigma_{jl}}{\sum_{k=1}^{m} \sigma_{jk}}$$

with $\sigma_{jk} \in \mathbf{R}^+$ we find

$$\prod_{j=1}^{n} \left(\frac{\sigma_{jl}}{\sum_{k=1}^{m} \sigma_{jk}}\right)^{\omega_j} \leq \sum_{j=1}^{n} \omega_j \frac{\sigma_{jl}}{\sum_{k=1}^{m} \sigma_{jk}}.$$

Summing over $1 \leq l \leq m$ gives

$$\sum_{l=1}^{m} \prod_{j=1}^{n} \left(\frac{\sigma_{jl}}{\sum_{k=1}^{m} \sigma_{jk}}\right)^{\omega_j} \leq 1$$

or

$$\sum_{l=1}^{m}\prod_{j=1}^{n}\sigma_{jl}^{\omega_j} \le \prod_{j=1}^{n}\left(\sum_{l=1}^{m}\sigma_{jl}\right)^{\omega_j}.$$

Setting $\sigma_{jk} = \rho_{jk}^{1/\omega_k}$,

$$\sum_{l=1}^{m}\prod_{j=1}^{n}\rho_{jl} \le \prod_{j=1}^{n}\left(\sum_{l=1}^{m}\rho_{jl}^{1/\omega_j}\right)^{\omega_j}.$$

If we now specialise to the case $n = 2$ and write $\xi_l = c_{1l}$, $\eta_l = c_{2l}$, $p = 1/\omega_1$ and $q = 1/\omega_2$ we find that

$$\sum_{l=1}^{m}\xi_l\eta_l \le \left(\sum_{l=1}^{m}\xi_l^p\right)^{1/p}\left(\sum_{l=1}^{m}\eta_l^q\right)^{1/q}$$

if

$$\frac{1}{p}+\frac{1}{q} = 1.$$

We have assumed here that $\xi_j, \eta_j > 0$, but this continues to hold if any are zero, since adding such terms to the sum leaves the left hand side unchanged and can only increase the right hand side. If we take $\chi_l = |\xi_l|$, $\psi_l = |\eta_l|$ with $\chi_l, \psi_l \in \mathbf{K}$ then we get

$$\sum_{l=1}^{m}|\chi_l\psi_l| \le \left(\sum_{l=1}^{m}|\chi_l|^p\right)^{1/p}\left(\sum_{l=1}^{m}|\psi_l|^q\right)^{1/q}.$$

Combining this with the triangle inequality gives *Hölder's inequality*

$$\left|\sum_{l=1}^{m}\chi_l\psi_l\right| \le \left(\sum_{l=1}^{m}|\chi_l|^p\right)^{1/p}\left(\sum_{l=1}^{m}|\psi_l|^q\right)^{1/q}.$$

The special case $p = q = 2$ of Hölder's inequality is known as the *Cauchy-Schwarz* inequality[1]

$$\left|\sum_{l=1}^{m}\chi_l\psi_l\right| \le \sqrt{\sum_{l=1}^{m}|\chi_l|^2}\sqrt{\sum_{l=1}^{m}|\psi_l|^2}.$$

[1] In the form presented in this section it is due to Cauchy. The corresponding integral inequality, presented in a later section, is due to Bunyakovsky, with a proof which was improved upon by Schwarz. The integral version and the version for infinite sums are also usually called Cauchy-Schwarz

Suppose $r \ge 1$ and $\tau_{jk} \in \mathbf{K}$ for $1 \le j \le m$, $1 \le k \le n$. Let

$$\upsilon_k = \sum_{j=1}^{m}\tau_{jk}.$$

From

$$|\upsilon_k|^r = |\upsilon_k||\upsilon_k|^{r-1} \le \sum_{j=1}^{m}|\tau_{jk}||\upsilon_k|^{r-1}$$

it follows that

$$\sum_{k=1}^{n}|\upsilon_k|^r \le \sum_{j=1}^{m}\sum_{k=1}^{n}|\tau_{jk}||\upsilon_k|^{r-1}.$$

Applying Hölder's inequality to the inner sum with $\xi_k = |\tau_{jk}|$, $\eta_k = |\upsilon_k|^{r-1}$, $p = r$, $q = r/(r-1)$, gives

$$\sum_{k=1}^{n}|\tau_{jk}||\upsilon_k|^{r-1} \le \left(\sum_{k=1}^{n}|\tau_{jk}|^r\right)^{1/r}\left(\sum_{k=1}^{n}|\upsilon_k|^r\right)^{(r-1)/r}$$

and hence

$$\sum_{k=1}^{n}|\upsilon_k|^r \le \left(\sum_{k=1}^{n}|\upsilon_k|^r\right)^{(r-1)/r}\sum_{j=1}^{m}\left(\sum_{k=1}^{n}|\tau_{jk}|^r\right)^{1/r}$$

and

$$\left(\sum_{k=1}^{n}|\upsilon_k|^r\right)^{1/r} \le \sum_{j=1}^{m}\left(\sum_{k=1}^{n}|\tau_{jk}|^r\right)^{1/r}.$$

This is *Minkowski's inequality.* It is easy to see that it holds also for $r = 1$. The special case $m = 2$, $\tau_{1k} = \xi_k$, $\tau_{2k} = \eta_k$ is also called Minkowski's inequality:

$$\left(\sum_{k=1}^{n}|\xi_k+\eta_k|^r\right)^{1/r} \le \left(\sum_{k=1}^{n}|\xi_k|^r\right)^{1/r} + \left(\sum_{k=1}^{n}|\eta_k|^r\right)^{1/r}.$$

1.2 Infinite Sums

Similar inequalities apply also with infinite sums in place of finite sums. These can be obtained just by taking limits. For example

$$\left|\sum_{l=1}^{m} \chi_l \psi_l\right| \leq \left(\sum_{l=1}^{m} |\chi_l|^p\right)^{1/p} \left(\sum_{l=1}^{m} |\psi_l|^q\right)^{1/q}.$$

for each m and, by definition,

$$\sum_{l=1}^{\infty} \chi_l \psi_l = \lim_{m\to\infty} \sum_{l=1}^{m} \chi_l \psi_l,$$

$$\sum_{l=1}^{\infty} |\chi_l|^p = \lim_{m\to\infty} \sum_{l=1}^{m} |\chi_l|^p,$$

and

$$\sum_{l=1}^{\infty} |\psi_l|^q = \lim_{m\to\infty} \sum_{l=1}^{m} |\psi_l|^q,$$

so

$$\left|\sum_{l=1}^{\infty} \chi_l \psi_l\right| \leq \left(\sum_{l=1}^{\infty} |\chi_l|^p\right)^{1/p} \left(\sum_{l=1}^{\infty} |\psi_l|^q\right)^{1/q}.$$

As we'll see, another way to obtain this is from the corresponding inequality for integrals, proved in the next section.

Unlike finite sums, infinite sums need not converge. The correct interpretation of this inequality is that the left hand side is finite if the right hand side is, in which case the left hand side is less than or equal to the right hand side. In this case that's a consequence of the fact that absolute convergence of a sum implies convergence, coupled with the fact that bounded monotone sequences converge. Similar remarks will apply to many inequalities in these notes, but from now on I will draw attention to them only if there is something particularly tricky in the proof of convergence.

A special case of Hölder's inequality is, as it was for finite sums, the Cauchy-Schwarz inequality.

$$\left|\sum_{l=1}^{\infty} \chi_l \psi_l\right| \leq \sqrt{\sum_{l=1}^{\infty} |\chi_l|^2}\sqrt{\sum_{l=1}^{\infty} |\psi_l|^2}.$$

The analogue of the Minkowski inequality for infinite sums is

$$\left(\sum_{k=1}^{\infty} |\xi_k + \eta_k|^r\right)^{1/r} \leq \left(\sum_{k=1}^{\infty} |\xi_k|^r\right)^{1/r} + \left(\sum_{k=1}^{\infty} |\eta_k|^r\right)^{1/r}.$$

1.3 Integrals

Similar inequalities apply also with integrals in place of sums. We could try to prove them by writing integrals as limits of finite sums, but it turns out to be easier just to mimic the arguments that led to the inequalities for finite sums. The two basic facts we need from the theory of integration are that if x is Lebesgue integrable and

$$x(t) \geq 0$$

for all t then

$$\int x(t) \geq 0$$

with strict inequality unless $x(t) = 0$ for almost all t, i.e. except for t in a set of Lebesgue measure zero, and that if x is Riemann integrable and

$$x(t) \geq 0$$

for all t then

$$\int x(t) \geq 0$$

with strict inequality unless $x(t) = 0$ at all points of continuity. The domain of integration will play no role in this section.

Recalling the inequality

$$\prod_{j=1}^{n} \theta_j^{\omega_j} \leq \sum_{j=1}^{n} \omega_j \theta_j$$

from the section on finite sums, we take $n = 2$, $\omega_1 = 1/p$ and $\omega_2 = 1/q$, where

$$\frac{1}{p} + \frac{1}{q} = 1.$$

This gives

$$\theta_1^{1/p}\theta_2^{1/q} \le \frac{\theta_1}{p} + \frac{\theta_2}{q}$$

with strict inequality unless $\theta_1 = \theta_2$. We apply this with

$$\theta_1 = \frac{|x(t)|^p}{\int |x(s)|^p\,ds}, \quad \theta_2 = \frac{|y(t)|^q}{\int |y(s)|^q\,ds}$$

where x and y are $\mathbf{K}$-valued measurable functions such that the integrals appearing in the denominator are finite and non-zero. This gives

$$\frac{|x(t)||y(t)|}{\left(\int |x(s)|^p\,ds\right)^{1/p}\left(\int |y(s)|^q\,ds\right)^{1/q}} \le \frac{1}{p}\frac{|x(t)|^p}{\int |x(s)|^p\,ds} + \frac{1}{q}\frac{|y(t)|^q}{\int |y(s)|^q\,ds}.$$

We then integrate this inequality over t to get

$$\frac{\int |x(t)||y(t)|\,dt}{\left(\int |x(s)|^p\,ds\right)^{1/p}\left(\int |y(s)|^q\,ds\right)^{1/q}} \le 1$$

or

$$\int |x(t)||y(t)|\,dt \le \left(\int |x(s)|^p\,ds\right)^{1/p}\left(\int |y(s)|^q\,ds\right)^{1/q}.$$

The absolute value of the integral is always bounded by the integral of the absolute value, so

$$\left|\int x(t)y(t)\,dt\right| \le \left(\int |x(s)|^p\,ds\right)^{1/p}\left(\int |y(s)|^q\,ds\right)^{1/q}.$$

which is Hölder's inequality for integrals.

If we define

$$x(t) = \chi_j, \quad y(t) = \psi_j$$

for $j-1 \le t < j$ then

$$\sum_{l=1}^{\infty} \chi_l\psi_l = \int_0^\infty x(t)y(t)\,dt,$$

$$\sum_{l=1}^{\infty} |\chi_l|^p = \int_0^\infty |x(t)|^p\,dt$$

and

$$\sum_{l=1}^{\infty} |\psi_l|^q = \int_0^\infty |y(t)|^q\,dt,$$

so

$$\left(\sum_{k=1}^{\infty} |\xi_k + \eta_k|^r\right)^{1/r} \le \left(\sum_{k=1}^{\infty} |\xi_k|^r\right)^{1/r} + \left(\sum_{k=1}^{\infty} |\eta_k|^r\right)^{1/r},$$

reproducing Hölder's inequality for infinite sums.

The derivation of Minkowski's inequality from Hölder's inequality follows the argument given earlier for finite sums, replacing the sums over k with integrals and keeping the sums over j as sums. Once again we can derive the inequality for infinite sums from the inequality for integrals by choosing x and y to be constant between non-negative integers.

As with infinite sums, we have to worry about convergence. The intended interpretation of integral inequalities where the left hand side is less than or equal to the right hand side is that if all functions involved are measurable and the right hand is finite then so is the left hand side and the inequality holds.

2 Normed Spaces

2.1 Metric Spaces

A *metric space* is a set E and a function function $d\colon E \times E \to \mathbf{R}$ with the properties

1. $d(x,y) \ge 0$ with $d(x,y) = 0$ if and only if $x = y$,
2. $d(x,y) = d(y,x)$,
3. $d(x,y) \le d(x,z) + d(y,z)$

for all $x, y, z \in E$. The inequality in the last of these is called the *triangle inequality* because in the case where E is the Euclidean plane and d is the usual notion of distance it says that the length of one side of a triangle is at most the sum of the lengths of the other two sides.

Before exploring the properties of metric spaces, here are some examples:

1. If $p \geq 1$ and $\mathbf{K}$ is either $\mathbf{R}$ or $\mathbf{C}$ then for any points $x = (\xi_1, \ldots, \xi_n)$, $y = (\eta_1, \ldots, \eta_n)$ in $E = \mathbf{K}^n$ we define
$$d(x, y) = \left(\sum_{\nu=1}^{n} |\xi_\nu - \eta_\nu|^p \right)^{1/p}.$$
This defines a metric space, which we will denote $l^p(n)$. The triangle inequality in this particular case is just Minkowski's inequality.

2. With $E = \mathbf{K}^n$ again, we can also take
$$d(x, y) = \max_{1 \leq \nu \leq n} |\xi_\nu - \eta_n|.$$
The resulting metric space is denoted $l^\infty(n)$.

3. The set $C([a, b])$ of continuous $\mathbf{K}$-valued functions on a closed interval $[a, b]$ with the metric
$$d(x, y) = \max_{a \leq t \leq b} |x(t) - y(t)|$$
is a metric space.

4. If T is a non-empty set then the set $B(T)$ of bounded $\mathbf{K}$-valued functions on T forms a metric space with metric
$$d(x, y) = \sup_{t \in T} |x(t) - y(t)|.$$
In the case $T = \{1, \ldots, n\}$ the supremum is a maximum, and we get the space $l^\infty(n)$ considered earlier.

5. The set (s) of all sequences in $\mathbf{K}$ with
$$d(x, y) = \sum_{n=1}^{\infty} \frac{1}{2^n} \frac{|\xi_n - \eta_n|}{1 + |\xi_n - \eta_n|}$$
is a metric space. The only part of this which is not obvious is the triangle inequality. To see this, note that if
$$0 \leq \sigma \leq \tau$$
for $\sigma, \tau \in \mathbf{R}$ then
$$\frac{\sigma}{1 + \sigma} \leq \frac{\tau}{1 + \tau}.$$
Apply this with $\sigma = |\alpha + \beta|$, $\tau = |\alpha| + |\beta|$ to obtain
$$\frac{|\alpha + \beta|}{1 + |\alpha + \beta|} \leq \frac{|\alpha| + |\beta|}{1 + |\alpha| + |\beta|} \leq \frac{|\alpha|}{1 + |\alpha|} + \frac{|\beta|}{1 + |\beta|},$$
Apply this with $\alpha = \xi_n - \zeta_n$ and $\beta = \zeta_n - \eta_n$ then multiply by $1/2^n$ and sum over n, to obtain
$$d(x, y) \leq d(x, z) + d(y, z).$$

6. Another example is the set E of functions from $\{1, \ldots, n\}$ to $\{0, 1\}$ and $d(x, y)$ the number of ν for which $x(\nu)$ and $y(\nu)$ differ. This d is known as the Hamming distance.

7. For *any* set E
$$d(x, y) = \begin{cases} 0, & \text{if } x = y \\ 1, & \text{if } x \neq y \end{cases},$$
is a metric, called the *discrete metric*.

We define *open* and *closed balls* in the usual way and then *open* and *closed subsets*.

On a metric space we can define a notion of convergence of sequences: $x_n \to y$ if and only if for all $\epsilon > 0$ there is an N such that $n > N$ implies $d(x_n, y) < \epsilon$.

For example, in $C([a, b])$ we have $x_n \to y$ if, for all $\epsilon > 0$ there is an N such that $n > N$ implies
$$\max_{a \leq t \leq b} |x_n(t) - y(t)| < \epsilon.$$
Equivalently, for all $\epsilon > 0$ there is an N such that for all $a \leq t \leq b$ and $n > N$ we have
$$|x_n(t) - y(t)| < \epsilon.$$
This is what is called *uniform convergence* in Analysis. It is not the same as *pointwise convergence*, which would mean that for all $\epsilon > 0$ and $a \leq t \leq b$ there is an N such that for all $n > N$ we have
$$|x_n(t) - y(t)| < \epsilon.$$

Uniform convergence implies pointwise convergence, but not conversely.

Returning to the general theory, it follows from the defining properties of a metric space that the limit, if it exists, is unique. If $x_n \to y$ then, for any $\epsilon > 0$, we have also $\epsilon/2 > 0$ and hence there is an N such that $n > N$ implies $d(x_n, y) < \epsilon/2$. Also, if $m > N$ then $d(x_m, y) < \epsilon/2$. It then follows from the triangle inequality that $d(x_m, x_n) < \epsilon$. We call a sequence *Cauchy* if, for every $\epsilon > 0$, there is an N such that $m, n > N$ implies $d(x_m, x_n) < \epsilon$, so what we've just shown is that every convergent sequence is Cauchy. A metric space in which the converse holds is called *complete*.

A subset of a metric space is a metric space. A complete subset is necessarily a closed subset. A closed subset of a complete metric space is also complete.

A subset of a metric space is called *compact* if every bounded sequence has a convergent subsequence.[2] Compact subsets of metric spaces are always closed and bounded. The converse is true in $l^p(n)$, but not in general.

Continuity of functions from one metric space to another can be defined either in the usual δ-ϵ way or by saying that the pre-image of an open set is open. These are equivalent. The image of a compact subset under a continuous function is compact. One easy consequence of this is that every continuous function from a compact set to $\mathbf{R}$ has both a minimum and a maximum.

2.2 Vector Spaces

A vector space E over $\mathbf{K}$ is defined via the usual axioms:

1. for every $x, y, z \in E$, $x + (y + z) = (x + y) + z$,
2. for every $x, y \in E$, $x + y = y + x$,
3. there is a $0 \in E$ such that for every $x \in E$, $x + 0 = x$,
4. for any $x \in E$ there is a $-x \in E$ such that $x + (-x) = 0$.
5. for every $\alpha \in \mathbf{K}$ and $x, y \in E$, $\alpha(x+y) = \alpha x+\alpha y$,
6. for every $\alpha, \beta \in \mathbf{K}$ and $x \in E$, $(\alpha + \beta)x = \alpha x + \beta x$,
7. for every $\alpha, \beta \in \mathbf{K}$ and $x \in E$, $(\alpha\beta)x = \alpha(\beta x)$, and
8. for every $x \in E$, $1x = x$.

The elementary consequences of these definitions, and the definition and properties of subspaces, sums, products, linearly independent and spanning sets will be assumed to be familiar, as will the fact that the space of linear functions on a vector space is also a vector space.

The following are examples of vector spaces:

1. the set $\mathbf{K}^n$ of n-tuples of elements of $\mathbf{K}$, including $\mathbf{K} = \mathbf{K}^1$,
2. the set (s) of all sequences,
3. the set l^p of all sequences $(\xi_1, \xi_2, \ldots)$ for which $\sum_{n\in\mathbf{Z}^+} |\xi_n|^p$ converges, for $1 \leq p < \infty$,
4. the set l^∞ of bounded sequences,
5. the set (c) of convergent sequences,
6. the set (c_0) of convergent sequences with limit 0,
7. the set $B(T)$ of bounded functions on a set T,
8. the set $C([a, b])$ of continuous functions on the closed bounded interval $[a, b]$,
9. the set $C_0(\mathbf{R})$ of continuous functions on $\mathbf{R}$ which vanish at infinity, i.e. for which
$$\lim_{t\to\infty} x(t) = 0,$$
10. The set $BV([a, b])$ of functions of bounded variation on $[a, b]$, i.e. those for which there is a uniform bound for $\sum_{1\leq j\leq n-1} |x(t_{j+1})-x(t_j)|$ for all increasing finite sequences $t_1 \leq \cdots \leq t_n \in [a, b]$.

[2] This definition turns out not to generalise well to more general topological spaces. A better definition from that point of view is that every open cover has a finite subcover. But the two statements are equivalent in the case of metric spaces and we'll find the sequential definition easier to use.

If e_k is the sequence whose k'th term is 1 and all others are zero then the set of all e_k is linearly independent in (s), l^p, l^∞, (c) and (c_0), but doesn't span any of these. Similarly the monomials t^k are linearly independent in $C([a,b])$ and $C^{(n)}([a,b])$ if $a < b$, but do not span.

2.3 Linear Transformations

Linear transformations are functions $A\colon E \to F$ from one $\mathbf{K}$-vector space to another such that

$$A(x+y) = Ax + Ay, \quad A(\alpha x) = \alpha Ax.$$

The set of such functions is itself a vector space and satisfies the identities

$$A(BC) = (AB)C, \quad A(B+C) = AB + AC,$$

$$(A+B)C = AB + BC, \quad \alpha(AB) = (\alpha A)B = A(\alpha B).$$

The identity transformation $I\colon E \to E$ given by $Ix = x$ serves as a multiplicative identity. Powers are defined in the usual way, including the inverse A^{-1}, if it exists. It satisfies the relation

$$(AB)^{-1} = B^{-1}A^{-1}.$$

The image and null space[3] are defined as usual. They are written as $A(E)$ and $N(E)$. Projections are linear transformations $P\colon E \to E$ such that $P^2 = P$.

2.4 Normed Spaces

A semi-norm on a vector space E is a function $p\colon E \to \mathbf{R}$ such that

1. for all $x \in E$, $p(x) \geq 0$,
2. for all $\alpha \in \mathbf{K}$ and $x \in E$, $p(\alpha x) = |\alpha| p(x)$, and
3. for all $x, y \in E$, $p(x,y) \leq p(x) + p(y)$.

[3]In Linear Algebra this is also often called the 'kernel', but this term is not used in Functional Analysis, because linear transformations are often defined as integral operators, and the word 'kernel' has a different meaning in that context.

A norm is a semi-norm such that $p(x) > 0$ for $x \neq 0$. Usually we write norms not using functional notation, as above, but as $\|x\|$, by analogy with absolute values. When there is ambiguity as to which of various possible norms is meant we may use subscripts, e.g. $\|x\|_p$.

The last inequality above, which becomes

$$\|x+y\| \leq \|x\| + \|y\|$$

in our new notation, is called the *triangle inequality*, because

$$d(x,y) = \|x - y\|$$

is a metric, called the *induced metric*, and the triangle inequality for norms is used in the proof of the corresponding inequality for metrics. In case the space is complete in the induced metric we call the normed space a *Banach space*.

A subspace of a normed space is also a normed space, but a subspace of a Banach space need not be a Banach space because, while a Cauchy sequence in the subspace is a Cauchy sequence also in the larger space, and hence convergent, its limit might not belong to the subspace. Closed subspaces of Banach spaces are however Banach spaces.

The following are examples of Banach spaces. Some have already been introduced as metric spaces and in each case the metric is easily seen to be the induced metric from the norm given below.

1. the space $l^p(n)$ with $1 \leq p < \infty$ and

   $$\|x\|_p = \left(\sum_{k=1}^{n} |\xi_k|^p\right)^{1/p}$$

   or $p = \infty$ and

   $$\|x\|_\infty = \max_{1 \leq k \leq n} |\xi_k|,$$

2. l^p with

   $$\|x\|_p = \left(\sum_{1 \leq k < \infty} |\xi_k|^p\right)^{1/p}$$

   for $1 \leq p < \infty$,

3. l^∞, the space of bounded sequences with

$$\|x\|_\infty = \sup_{k=1}^{\infty} |\xi_k|,$$

4. the space (c) of convergent sequences, with the same norm as l^∞, of which it is a subset,

5. the space (c_0) of sequences converging to 0, again with the same norm,

6. the space $B(T)$ of bounded functions on a set T, with the norm

$$\|x\| = \sup_{t\in T} |x(t)|,$$

7. the subspace $C([a,b])$ of $B([a,b])$, consisting of bounded continuous functions with the same norm.

The completeness of l^p is not obvious. Suppose $(x_1, x_2, \ldots)$ is a Cauchy sequence, so for any $\epsilon' > 0$ there is an N such that for $k, l > N$,

$$\|x_k - x_l\|_p = \left(\sum_{1\le n<\infty} |\xi_n^{(k)} - \xi_n^{(l)}|^p \right)^{1/p} < \epsilon'.$$

Here

$$x_k = (\xi_1^{(k)}, \xi_2^{(k)}, \ldots).$$

This notation isn't great, but when we're dealing with sequences of sequences there really isn't a notation which isn't at least somewhat confusing. Then for each n we have

$$|\xi_n^{(k)} - \xi_n^{(l)}| < \epsilon'$$

and so $(\xi_n^{(1)}, \xi_n^{(2)}, \ldots)$ is a Cauchy sequence in $\mathbf{K}$, and hence converges. Call its limit η_n. From

$$\left(\sum_{n=1}^{\infty} |\xi_n^{(k)} - \xi_n^{(l)}|^p \right)^{1/p} < \epsilon'$$

it follows that

$$\left(\sum_{n=1}^{m} |\xi_n^{(k)} - \xi_n^{(l)}|^p \right)^{1/p} < \epsilon'$$

for each $m \in \mathbf{Z}+$. We can exchange limits with differences, absolute values, powers and finite sums. We can also exchange limits with non-strict inequalities so, taking the limit as $l \to \infty$,

$$\left(\sum_{n=1}^{m} |\xi_n^{(k)} - \eta_n|^p \right)^{1/p} \le \epsilon'.$$

Then, taking the limit as $m \to \infty$,

$$\left(\sum_{n=1}^{\infty} |\xi_n^{(k)} - \eta_n|^p \right)^{1/p} \le \epsilon'$$

and so the sequence $(\xi_1^{(k)} - \eta_1, \xi_2^{(k)} - \eta_2, \ldots)$ belongs to l^p and then, because l^p is a vector space, and hence closed under addition, $y = (\eta_1, \eta_2, \ldots)$ belongs to l^p. Thus

$$\|x_k - y\| \le \epsilon'$$

for $k > N$. Now, for any $\epsilon > 0$ we can find an $\epsilon' > 0$ with $\epsilon' < \epsilon$ and hence, with N as above,

$$\|x_k - y\| < \epsilon$$

for $k > N$. Since ϵ was arbitrary except for the requirement that $\epsilon > 0$ it follows that $x_k \to y$ in l^p.

The completeness of (c) is even farther from obvious. Suppose $x_j = (\xi_1^{(j)}, \xi_2^{(j)}, \ldots) \in (c)$ and x_j is a Cauchy sequence in (c). By definition, for each $\epsilon' > 0$ there is an $M \in \mathbf{Z}^+$ such that $j, k > M$ implies

$$\|x_j - x_k\| = \sup_{l\in\mathbf{Z}^+} |\xi_l^{(j)} - \xi_l^{(k)}| < \epsilon',$$

i.e. that

$$|\xi_l^{(j)} - \xi_l^{(k)}| < \epsilon'$$

for all $j, k > M$ and $l \in \mathbf{Z}^+$. It follows that for each $l \in \mathbf{Z}^+$, $(\xi_l^{(1)}, \xi_l^{(2)}, \ldots)$ is a Cauchy sequence in $\mathbf{K}$ and hence has a limit

$$\eta_l = \lim_{j\to\infty} \xi_l^{(j)}.$$

Also, taking the limit $k \to \infty$ in the inequality above,

$$|\xi_l^{(j)} - \eta_l| \le \epsilon'$$

for all $j > M$ and $l \in \mathbf{Z}^+$. Note that this M depends only on ϵ'.

Now $x_j \in (c)$ so there is a $\zeta^{(j)}$ such that

$$\lim_{l\to\infty} \xi_l^{(j)} = \zeta^{(j)}.$$

In other words, for each $\epsilon' > 0$ there is an N_j such for $l > N_j$ we have

$$|\xi_l^{(j)} - \zeta^{(j)}| < \epsilon'.$$

It follows that $z = (\zeta^{(1)}, \zeta^{(2)}, \cdots)$ is a Cauchy sequence. Indeed, for any $\epsilon > 0$ we can take[4]

$$\epsilon' = \frac{\epsilon}{6}$$

Clearly $\epsilon' > 0$. There is therefore an M, as above such that

$$|\xi_l^{(j)} - \xi_l^{(k)}| < \epsilon'$$

for all $l > M$ and $j, k \in \mathbf{Z}^+$. This applies in particular if $l > \max(N_j, N_k)$, in which case

$$|\xi_l^{(j)} - \zeta^{(j)}| < \epsilon', \quad |\xi_l^{(k)} - \zeta^{(k)}| < \epsilon'.$$

Then, by the triangle inequality,

$$|\zeta^{(j)} - \zeta^{(k)}| < \epsilon' + \epsilon' + \epsilon' < \epsilon.$$

Thus there is, for each $\epsilon > 0$, an M such that if $j, k > M$ then

$$|\zeta^{(j)} - \zeta^{(k)}| < \epsilon,$$

i.e. z is a Cauchy sequence in $\mathbf{K}$. It is therefore convergent. Define

$$\omega = \lim_{j\to\infty} \zeta^{(j)}.$$

Taking the limit in the inequality

$$|\zeta^{(j)} - \zeta^{(k)}| < \epsilon' + \epsilon' + \epsilon'$$

gives

$$|\zeta^{(j)} - \omega| \leq 3\epsilon',$$

which will be useful later.

We next show that

$$\lim_{l\to\infty} \eta_l = \omega.$$

There is, as we've already established, an $M > 0$ such that

$$|\xi_l^{(j)} - \eta_l| \leq \epsilon'$$

for $j > M$ and $l \in \mathbf{Z}^+$ and

$$|\zeta^{(j)} - \omega| \leq 3\epsilon'$$

for $j > M$. There is also an N_j such that

$$|\xi_l^{(j)} - \zeta^{(j)}| < \epsilon'$$

for $l > N_j$. Choose any such j and then any such l. Then by the triangle inequality

$$|\eta_l - \omega| < \epsilon' + 3\epsilon' + \epsilon' = 5\epsilon' < \epsilon.$$

So there is for every $\epsilon > 0$ an M such that

$$|\eta_l - \omega| < \epsilon$$

for all $l > M$. In other words,

$$\lim_{l\to\infty} \eta_l = \omega.$$

This shows that

$$y \in (c).$$

Finally, we show that

$$\lim_{j\to\infty} x_j = y$$

in the norm on (c). We've already seen that there is, for each $\epsilon > 0$, an M such that for all $j > M$

$$|\xi_l^{(j)} - \eta_l| \leq 5\epsilon'$$

and hence

$$\|x_j - y\| = \sup_{l\in\mathbf{Z}^+} |\xi_l^{(j)} - \eta_l| \leq 5\epsilon' < \epsilon.$$

So for any $\epsilon > 0$ there is a $j > M$ such that

$$\|x_j - y\| \leq \epsilon < \epsilon''',$$

[4]Clearly the 6 in the equation below could be replaced by 3, but we will leave some slack for later parts of the argument

and we're done.

The fact that (c_0) is complete is just the above argument, plus the observation that the relation

$$\lim_{l\to\infty} \eta_l = \omega,$$

proved above, implies that if

$$\zeta^{(j)} = \lim_{l\to\infty} \xi_l^{(j)} = 0$$

for all $j \in \mathbf{Z}^+$ then

$$\lim_{l\to\infty} \eta_l = \omega = \lim_{j\to\infty} \zeta^{(j)} = 0.$$

2.5 Bounded Linear Transformations

Suppose E and F are normed spaces and $A\colon E \to F$ is a linear transformation A $\lambda \in \mathbf{R}$ such that for all $x \in E$ we have

$$\|Ax\| \le \lambda \|x\|,$$

is called a *bound* for A. There need not be any bound. Note that the norms on the different sides of the inequality are generally different. On the left we have the norm on F, while on the right we have the norm on E. Linear transformations which have a bound are called *bounded*. The set of all such transformations is written $\mathcal{L}(E, F)$, with $\mathcal{L}(E, E)$ being abbreviated to $\mathcal{L}(E)$.

If $A \in \mathcal{L}(E, F)$ and $x_n \to x$ in E then

$$\|Ax_n - Ax\| = \|A(x_n - x)\| \le \lambda \|x_n - x\| \to 0,$$

from which it follows that $Ax_n \to Ax$. In other words, A is continuous. On the other hand, if $A\colon E \to F$ is linear but not bounded then no $n \in \mathbf{Z}^+$ is a suitable choice of λ, so there is an $x_n \in E$ such that

$$\|Ax_n\| > n\|x_n\|.$$

Then, setting

$$y_n = \frac{x_n}{\|x_n\|}, \quad z_n = \frac{y_n}{\|Ay_n\|} = \frac{x_n}{\|Ax_n\|},$$

we have

$$\|y_n\| = 1, \quad \|Ay_n\| \to \infty,$$
$$z_n \to 0, \quad \|Az_n\| = 1.$$

This is incompatible with $Az_n \to 0$, so A is not continuous. In other words, linear transformations are bounded if and only if they are continuous.

If $E \neq 0$ then any bound is non-negative, as we will show below, so the set of bounds, if it's non-empty has an infimum. In fact it must have a minimum, since we can, for each $x \in E$, take the limit as λ tends to its infimum in the inequality

$$\frac{\|Ax\|}{\|x\|} \le \lambda.$$

The minimum bound is called the norm of A and is denoted $\|A\|$. To justify this notation and terminology we need to check that this satisfies the defining properties of a norm on the vector space $\mathcal{L}(E, F)$.

First we check that $\|A\| \ge 0$ with equality if and only if $A = 0$, assuming still that $E \neq 0$. If $\|A\| < 0$ and $x \neq 0$ then $\|x\| > 0$ and $\|Ax\| \le \|A\|\|x\| < 0$, which is impossible. If $A = 0$ then 0 is clearly a bound and all bounds are non-negative, so it must be the minimal bound $\|A\|$. Conversely, if $\|A\| = 0$ then $\|Ax\| \le \|A\|\|x\| = 0$ for all $x \in E$ and hence $Ax = 0$ for all $x \in E$ and therefore $A = 0$.

Next we show that $\|\alpha A\| = |\alpha|\|A\|$. We observe that for any $x \in E$

$$\|\alpha Ax\| = |\alpha|\|Ax\| \le |\alpha|\|A\|\|x\|.$$

From this it follows that $|\alpha|\|A\|$ is a bound for αA, but we don't yet know that it's the minimal bound. In other words, all we know at this stage is that

$$\|\alpha A\| \le |\alpha|\|A\|.$$

It's clear that

$$\|\alpha A\| = |\alpha|\|A\|$$

if $\alpha = 0$, because both sides of the equation are then zero. If $\alpha \neq 0$ then we observe that, for any $x \in E$,

$$\|\alpha^{-1}\alpha Ax\| = |\alpha^{-1}|\|\alpha Ax\| \le |\alpha^{-1}|\|\alpha A\|\|x\|.$$

From this it follows that $|\alpha^{-1}|\|\alpha A\|$ is a bound for $\alpha^{-1}\alpha A = A$, so

$$\|A\| \le |\alpha^{-1}|\|\alpha A\| = |\alpha|^{-1}\|\alpha A\|$$

or, equivalently,

$$\|\alpha A\| \geq |\alpha| \|A\|.$$

This combined with the reverse inequality, which we derived previously, implies

$$\|\alpha A\| = |\alpha| \|A\|.$$

Finally we show that $\|A + B\| \leq \|A\| + \|B\|$. For any $x \in E$,

$$\begin{aligned} \|(A + B)x\| &= \|Ax + Bx\| \leq \|Ax\| + \|Bx\| \\ &\leq \|A\|\|x\| + \|B\|\|x\| \\ &= (\|A\| + \|B\|)\|x\|. \end{aligned}$$

Thus $\|A\| + \|B\|$ is a bound for $A + B$. The minimal bound is at most this large, so

$$\|A + B\| \leq \|A\| + \|B\|.$$

If $A \in (F, G)$ and $B \in (E, F)$ then, for all $x \in E$,

$$\|ABx\| \leq \|A\|\|Bx\| \leq \|A\|\|B\|\|x\|.$$

Thus $\|A\|\|B\|$ is a bound for AB, which means that $AB \in \mathcal{L}(E, G)$ and

$$\|AB\| \leq \|A\|\|B\|.$$

From the preceding result we can prove by induction on n that

$$\|A^n\| \leq \|A\|^n$$

if $A \in \mathcal{L}(E)$. The inequality can be strict, as we can see by taking E to be finite dimensional and A to be nilpotent, but non-zero.

If $A_n \to A$ in $\mathcal{L}(E, F)$ then $A_n x \to Ax$ in F for all $x \in E$. This is easily seen, since

$$\|A_n x - Ax\| = \|(A_n - A)x\| \leq \|A_n - A\|\|x\| \to 0$$

if $A_n \to A$. The converse is however false. For a counterexample, consider $A_n \in \mathcal{L}(l^1)$ defined by

$$A_n(\xi_1, \xi_2, \ldots) = (\xi_1, \ldots, \xi_n, 0, 0, \ldots).$$

If F is a Banach space then so is $\mathcal{L}(E, F)$. To see this, suppose $(A_1, A_2, \ldots)$ is a Cauchy sequence in $\mathcal{L}(E, F)$, i.e. that there is for every $\epsilon > 0$ an N such that if $j, k > N$ then

$$\|A_j - A_k\| < \epsilon.$$

If $x \in E$ is non-zero and $\epsilon' > 0$ then

$$\epsilon = \frac{\epsilon'}{\|x\|} > 0$$

and

$$\begin{aligned} \|A_j x - A_k x\| = \|(A_j - A_k)x\| &\leq \|A_j - A_k\|\|x\| \\ &< \epsilon\|x\| = \epsilon'. \end{aligned}$$

So there is, for each $\epsilon' > 0$, an N such that $j, k > N$ implies

$$\|A_j x - A_k x\| < \epsilon'.$$

In other words, $(A_1 x, A_2 x, \ldots)$ is a Cauchy sequence in F. But F is complete, so the sequence has a limit, which we call Ax. We also define $A0 = 0$. It is straightforward to show that A, thus defined, is linear and bounded, and that $A_n \to A$ in $\mathcal{L}(E, F)$.

Addition, scalar multiplication and multiplication are all continuous. If $A_n \to A$ and $B_n \to B$ then $A_n + B_n \to A + B$. Given $\epsilon > 0$ set $\epsilon' = \epsilon/2$. Then there are N_A and N_B such that $j > N_A$ implies

$$\|A_j - A\| < \epsilon'$$

and $j > N_B$ implies

$$\|B_j - B\| < \epsilon'.$$

If $N = \max(N_A, N_B)$ then $j > N$ implies

$$\begin{aligned} \|(A_j + B_j) - (A + B)\| &= \|(A_j - A) + (B_j - B)\| \\ &\leq \|(A_j - A)\| \\ &\quad + \|(B_j - B)\| \\ &< \epsilon' + \epsilon' = \epsilon. \end{aligned}$$

So $A_n + B_n \to A + B$, as promised. This proof is the same as the proof of the corresponding property of real or complex numbers. The proofs for multiplication also follow from those for real or complex numbers, just replacing absolute values by norms where needed.

One thing which may be unexpected if your intuition is based on finite-dimensional vector spaces is that there are $A \in \mathcal{L}(E)$ with zero nullspace which are not invertible. An example in $(C[0,1])$ is

$$(Ax)(t) = tx(t).$$

We have $|tx(t)| \leq |x(t)|$ for $t \in [0,1]$, so

$$\begin{aligned} \|Ax\| &= \max_{0\leq t\leq 1} |(Ax)(t)| = \max_{0\leq t\leq 1} |tx(t)| \\ &\leq \max_{0\leq t\leq 1} |x(t)| = \|x\|. \end{aligned}$$

Thus A is bounded.

If $Ax = 0$ then $tx(t) = 0$ for $0 \leq t \leq 1$ so $x(t) = 0$ for $0 < t \leq 1$ and then, by continuity, for $0 \leq t \leq 1$, so $x = 0$. So the nullspace of A is 0.

But A is not invertible. There is no $x \in C([0,1])$ such that $Ax = y$ where

$$y(t) = \sqrt{t}.$$

2.6 Normed Spaces of Finite Dimension

Suppose $x_1, \ldots x_n$ are linearly independent vectors in a normed vector space E. Then there is a $\mu \in \mathrm{R}^+$ such that

$$\sum_{j=1}^{n} |\beta_j| \leq \mu \left\| \sum_{j=1}^{n} \beta_j x_j \right\|$$

for all $\beta_1, \ldots, \beta_n \in \mathbf{K}$. To see this, consider $X\colon l^1(n) \to E$ defined by

$$X((\alpha_1, \ldots, \alpha_n)) = \alpha_1 x_1 + \cdots + \alpha_n x_n.$$

Let $a = (\alpha_1, \ldots, \alpha_n) \in l^1(n)$. By the properties of the norm,

$$\begin{aligned} \|Xa\| &= \|\alpha_1 x_1 + \cdots + \alpha_n x_n\| \\ &\leq |\alpha_1|\|x_1\| + \cdots + |\alpha_n|\|x_n\| \\ &\leq \lambda(|\alpha_1| + \ldots |\alpha_n|) = \lambda\|a\|, \end{aligned}$$

where

$$\lambda = \max_{1\leq j\leq n} \|x_j\|.$$

So X is bounded and hence continuous. The set

$$M = \{a \in l^1(n) \colon \|a\| = 1\}$$

is a closed and bounded subset of the finite dimensional space $l^1(n)$ and thus is compact. So then is XM. The norm is a continuous function on XM and hence has a minimum, which we will call μ. It must be non-negative, but it can't be 0, since $x_1, \ldots, x_n$ are linearly independent. It follows that $\mu > 0$. If $\beta_1, \ldots, \beta_n$ are not all zero then, defining

$$\alpha_j = \beta_j / \sum_{k=1}^{n} |\beta_k|,$$

we have

$$\sum_{j=1}^{n} |\alpha_j| = 1,$$

which, by the definition of μ, implies

$$1 \leq \mu \left\| \sum_{j=1}^{m} \alpha_j x_j \right\|.$$

Equivalently,

$$\sum_{j=1}^{n} |\beta_j| \leq \mu \left\| \sum_{j=1}^{m} \beta_j x_j \right\|.$$

The same is clearly still true when we drop the assumption that $\beta_1, \ldots, \beta_n$ are not all zero. This is exactly what was claimed earlier.

It follows from this that if $\{x_1, \ldots, x_n\}$ is a basis for E,

$$y_k = \sum_{j=1}^{n} \alpha_j^{(k)} x_j,$$

and

$$z = \sum_{j=1}^{n} \gamma_j x_j$$

then $y_n \to z$ if and only if $\alpha_j^{(n)} \to \gamma_j$ for each j. Since this condition is independent of the choice of norm it follows that all norms on a finite dimensional vector space are equivalent. In particular, every finite dimensional vector space is complete, and therefore also closed. A further consequence is that every linear function from finite dimensional vector space to a a normed vector space is continuous. To see this, suppose $A\colon E \to F$ is linear and that $\{x_1, \ldots, x_n\}$ is a basis for E. Suppose further that $y_n \to z$ in E. Let the coefficients with respect to the given basis be

$$y_k = \sum_{j=1}^{n} \alpha_j^{(k)} x_j,$$

and

$$z = \sum_{j=1}^{n} \gamma_j x_j.$$

Then, as we saw above, $\alpha_j^{(n)} \to \gamma_j$ for each j. Then

$$\begin{aligned} Ay_k = \sum_{j=1}^{n} \alpha_j^{(k)} Ax_j &\to \sum_{j=1}^{n} \gamma_j Ax_j \\ &= A\left(\sum_{j=1}^{n} \gamma_j x_j\right) = Az. \end{aligned}$$

The *Riesz Lemma* states that if F is a closed proper subspace of a normed space E then there is, for each $\eta \in (0,1)$, a vector $x \in E$ such that

$$\|x\| = 1$$

and, for all $y \in F$,

$$\|x - y\| \geq \eta.$$

To prove this, first note that F is a proper subspace, so there is a $z \notin F$. Set

$$\delta = \inf_{y \in F} \|y - z\|.$$

There is then a sequence y_n such that $\|y_n - z\| \to \delta$. We know that $\delta > 0$, because otherwise $y_n \to z$, which would contradict the assumption that F is closed. Now $\eta \in (0,1)$, so $\delta/\eta > \delta$ and hence there is a $w \in F$ with $0 < \|w - z\| \leq \delta/\eta$. Let $\gamma = 1/\|w - z\|$ and

$$x = -\gamma(w - z).$$

Then $\|x\| = 1$ and, for all $y \in F$,

$$\begin{aligned} \|y - x\| &= \|y + \gamma(w - z)\| = \|(y + \gamma w - \gamma z\| \\ &= \gamma \left\| \left(\frac{1}{\gamma} y + w\right) - z \right\| \geq \frac{\eta}{\delta}\delta = \eta. \end{aligned}$$

since $\gamma \geq \eta/\delta$ and $(1/\gamma)y + w \in F$.

As a corollary, a normed vector space E is finite dimensional if and only if every bounded sequence has a convergent subsequence.

Suppose that E is finite dimensional, i.e. that $\{x_1, \ldots, x_n\}$ is a basis for E, and $(y_1, y_2, \ldots)$ is a bounded sequence in E, i.e. that there is a γ such that $\|y_k\| \leq \gamma$ for all $k \in \mathbf{Z}^+$. Take coordinates with respect to the basis,

$$y_k = \sum_{j=1}^{n} \alpha_j^{(k)} x_j.$$

Then there is a μ such that

$$\sum_{j=1}^{n} |\alpha_j^{(k)}| \leq \mu \left\| \sum_{j=1}^{n} \alpha_j^{(k)} x_j \right\| = \mu \|y_k\| \leq \mu\gamma.$$

Let

$$a_k = (\alpha_1^{(k)}, \ldots, \alpha_n^{(k)}).$$

Then $(a_1, a_2, \ldots)$ is a bounded sequence in $l^1(n)$. It therefore has a convergent subsequence $a_{k_l} \to b = (\beta_1, \ldots, \beta_n)$. But then, for each $j \in \{1, \ldots, n\}$, we have $\alpha_j^{(k_l)} \to \beta_j$ as $l \to \infty$. From this it follows, as we've already seen, that $y_{k_l} \to z$, where $z = \beta_1 x_1 + \cdots + \beta_n x_n$. So if E is finite dimensional then every bounded sequence has a convergent subsequence.

Suppose now that E is infinite dimensional and choose $x_1 \in E$ with

$$\|x_1\| = 1.$$

Let F_1 be the span of $\{x_1\}$. Being finite dimensional, it must be closed and so, by the Riesz Lemma with $\eta = 1/2$, there is an $x_2 \in E$ with

$$\|x_2\| = 1$$

and

$$\|x_2 - y\| \geq 1/2$$

for all $y \in F_1$. Let F_2 be the span of $\{x_1, x_2\}$ and find an x_3 with

$$\|x_3\| = 1$$

and

$$\|x_3 - y\| \geq 1/2$$

for all $y \in F_2$. Since E is infinite dimensional we can continue this procedure indefinitely, constructing a sequence $(x_1, x_2, \ldots)$ in E which is bounded, because

$$\|x_k\| = 1$$

for all k, but no subsequence of which is convergent, because

$$\|x_{k_l} - x_{k_m}\| \geq 1/2$$

for all $l, m \in \mathbf{Z}^+$ and so the Cauchy criterion fails for $\epsilon = 1/2$. So if E is infinite dimensional then there is a bounded sequence with no convergent subsequence.

It follows from this, and from the definition of compactness, that the following statements are equivalent for a normed space E:

1. E is finite dimensional.
2. Every closed bounded subset of E is compact.
3. The closed unit ball in E is compact.

2.7 Geometric Series

Convergence of series in a normed space is defined in the usual way, as a limit of finite sums:

$$\sum_{j=i}^{\infty} x_j = \lim_{n\to\infty} \sum_{j=1}^{n} x_j .$$

If

$$\sum_{j=i}^{\infty} \|x_j\|$$

converges then

$$\sum_{j=i}^{\infty} x_j$$

is Cauchy. This follows immediately from

$$\left\| \sum_{j=i}^{k} x_j - \sum_{j=i}^{l} x_j \right\| \leq \sum_{j=i}^{k} \|x_j\| - \sum_{j=i}^{l} \|x_j\|,$$

which is a consequence of the triangle inequality, combined with the fact that

$$\sigma_n = \sum_{j=i}^{n} \|x_j\|.$$

is a Cauchy sequence. If our normed space is a Banach space then

$$\sum_{j=i}^{\infty} x_j$$

converges. It is easy to see then that

$$\left\| \sum_{j=i}^{\infty} x_j \right\| \leq \sum_{j=i}^{\infty} \|x_j\|.$$

Indeed the partial sums all have the right hand side as a bound, as a consequence of the triangle inequality, so the infinite sum must as well, since it's a limit of partial sums and the norm is continuous. So the comparison test works for series in a Banach space.

If $K \subseteq \mathcal{L}(E)$ and

$$\sum_{j=0}^{\infty} K^j$$

converges then, using the fact that multiplication is continuous,

$$K \sum_{j=0}^{\infty} K^j = \sum_{j=1}^{\infty} K^j = \sum_{j=0}^{\infty} K^j - I$$

and similarly

$$\left(\sum_{j=0}^{\infty} K^j \right) K = \sum_{j=0}^{\infty} K^j - I$$

and hence

$$\sum_{j=0}^{\infty} K^j = (I - K)^{-1} .$$

The series need not converge, of course, but it will, by the comparison test, if

$$\|K\| < 1$$

and the space is a Banach space. In fact, we have

$$\left\| (I - K)^{-1} \right\| \leq (1 - \|K\|)^{-1}$$

in that case.

2.8 Normed Algebras

An algebra is a vector space with a multiplication satisfying

$$x(yz) = (xy)z, \quad \alpha(xy) = (\alpha x)y = x(\alpha y),$$

$$x(y+z) = xy + xz, \quad (x+y)z = xz + yz.$$

A commutative algebra satisfies

$$xy = yx$$

in addition. Since we're considering vector spaces there is automatically a zero element, i.e. an additive identity. There may or may not be a multiplicative identity, i.e. an element e such that

$$xe = x = ex.$$

This element, if it exists, is necessarily unique.

For example, the space $\mathcal{L}(E)$, where E is a normed space, is a non-commutative[5] algebra with the identity transformation as multiplicative identity. The space $C([a,b])$, with multiplication defined by

$$(xy)(t) = x(t)y(t)$$

is a commutative algebra with a multiplicative identity consisting of the function which is equal to 1 everywhere.

A *normed algebra* is a normed space which is also an algebra with multiplication satisfying

$$\|xy\| \leq \|x\|\|y\|.$$

As in the special case of $\mathcal{L}(E)$, addition, scalar multiplication and multiplication are necessarily continuous and

$$\|x^n\| \leq \|x\|^n$$

in any normed algebra.

A normed algebra whose induced metric is complete is called a *Banach algebra*. An example of a commutative Banach algebra is the space $l^1(\mathbf{Z})$ consisting of functions on $x\colon \mathbf{Z} \to \mathbf{K}$ with

$$\sum_{j\in\mathbf{Z}} |x(j)| < \infty.$$

[5]provided $\dim(E) > 1$

Addition, scalar multiplication and multiplication are defined by

$$(x+y)(j) = x(j) + y(j), \quad (\alpha x)(j) = \alpha x(j),$$

$$(xy)(j) = \sum_{k\in\mathbf{Z}}^{\infty} x(j-k)y(k).$$

It has the multiplicative identity

$$e(j) = \begin{cases} 1 & \text{if } j = 0 \ , \\ 0 & \text{if } j \neq 0 \ . \end{cases}$$

Just as for sequences, i.e. functions from $\mathbf{Z}^+$ to $\mathbf{K}$, we often use subscript notation in place of functional notation. In other words, we write ξ_j in place of $x(j)$.

Geometric series may be defined in a Banach algebra with multiplicative identity, just as they were in $\mathcal{L}(E)$. If $\|x\| < 1$ in such an algebra then $e - x$ is invertible, with

$$(e-x)^{-1} = \sum_{j=0}^{\infty} x^j.$$

More generally, if y is invertible then so is x for all x with $\|x - y\| < 1/\|y^{-1}\|$. To see this, observe that

$$\|y^{-1}(y-x)\| \leq \|y^{-1}\|\|y-x\| < 1$$

and so $e - y^{-1}(y-x) = y^{-1}x$ is invertible. Let

$$z = \left(y^{-1}x\right)^{-1} y^{-1}$$

Then $zx = e$. Similarly

$$\|(y-x)y^{-1}\| \leq \|y-x\|\|y^{-1}\| < 1$$

so $e - (y-x)y^{-1} = xy^{-1}$ is invertible. Let

$$w = y^{-1}\left(xy^{-1}\right)^{-1}.$$

Then $xw = e$. It follows that

$$w = ew = zxw = ze = z$$

and so

$$zx = e = xz.$$

In other words, z is an inverse to x, which is therefore invertible. Since each invertible element is contained in a ball where every element is invertible we conclude that the set of invertible elements in a Banach algebra is open.

With x and y as above we have

$$x^{-1} - y^{-1} = y^{-1}(y-x)\left(e - y^{-1}(y-x)\right)^{-1} y^{-1},$$

from which it follows that

$$\left\|x^{-1} - y^{-1}\right\| \le \frac{\left\|y^{-1}\right\|^2 \|x-y\|}{1 - \|y^{-1}\| \, \|x-y\|}.$$

From this it follows that inversion is a continuous function on the set of invertible elements.

2.9 The Banach Fixed Point Theorem

Suppose that E is a complete metric space with metric d and $X \subseteq E$. A function $f\colon X \to E$ is said to be a *contraction mapping* on X if there is a $q \in (0,1)$ such that for all $x, y \in X$

$$d(f(x), f(y)) \le q d(x,y).$$

This implies in particular that f is continuous on X.

Given $x_0 \in X$ we can define a sequence $(x_1, x_2, \ldots)$ i E by

$$x_{n+1} = f(x_n)$$

for $n \ge 0$. If the ball of radius

$$\rho = \frac{1}{1-q} d(x_0, f(x_0))$$

about x_0 is contained in X then, by generalised induction,

$$d(x_n, x_{n+1}) \le q^n d(x_0, f(x_0)),$$

$$d(x_0, x_{n+1}) \le \frac{1-q^{n+1}}{1-q} d(x_0, f(x_0))$$

$$x_{n+1} \in X.$$

The first two statements are clear for $n = 0$, because they then reduce to the $d(x_0, x_1)$ is less than or equal to itself. The third of these statements follows, not just for $n = 0$ but for all $n \ge 0$, from the second, because

$$\frac{1-q^{n+1}}{1-q} d(x_0, f(x_0)) < \rho.$$

Suppose then the statements above hold for $n = k$. Then $x_k, x_{k+1} \in X$, so

$$\begin{aligned} d(x_{k+1}, x_{k+2}) &= d(f(x_k), f(x_{k+1}) \le q d(x_k, x_{k+1}) \\ &= \le q q^k d(x_0, f(x_0)) \\ &= q^{k+1} d(x_0, f(x_0)). \end{aligned}$$

Then

$$\begin{aligned} d(x_0, x_{k+2}) &\le d(x_0, x_{k+1}) + d(x_{k+1}, x_{k+2}) \\ &\le \frac{1-q^{k+1}}{1-q} d(x_0, f(x_0)) \\ &\qquad + q^{k+1} d(x_0, f(x_0)) \\ &= \frac{1-q^{k+2}}{1-q} d(x_0, f(x_0)). \end{aligned}$$

These are the first two of our three statements with $n = k+1$ and the third, as noted previously, follows from these, so this completes the induction. A further induction shows that if $j \le k$ then

$$d(x_j, x_k) \le \frac{q^j - q^{k+1}}{1-q} d(x_0, f(x_0)).$$

It then follows that if $j, k > N$ then

$$d(x_j, x_k) < \frac{q^{N+1}}{1-q} d(x_0, f(x_0)).$$

so there is, for every $\epsilon > 0$, an N such that $j, k > N$ implies

$$d(x_j, x_k) < \epsilon.$$

In fact any N greater than or equal to

$$1 + \frac{\log\left(\frac{(1-q)\epsilon}{d(x_0, f(x_0))}\right)}{\log q}$$

will work if $f(x_0) \ne x_0$ and any N will work if $x_0 = f(x_0)$. So the sequence $(x_1, x_2, \ldots)$ is Cauchy, and hence converges to a limit point in X. Let

$$z = \lim_{j\to\infty} x_j.$$

Then

$$f(z) = \lim_{j\to\infty} f(x_j) = \lim_{j\to\infty} x_{j+1} = z.$$

So z is a fixed point of f. Taking limits in

$$d(x_0, x_{n+1}) \leq \frac{1 - q^{n+1}}{1 - q} d(x_0, f(x_0))$$

shows that

$$d(x_0, z) \leq \rho.$$

Suppose $w \in X$ is a fixed point of f, i.e. that

$$w = f(w).$$

Then

$$d(w, z) = d(f(w), f(z)) \leq q d(w, z)$$

and hence

$$(1 - q) d(w, z) \leq 0.$$

But $1 - q > 0$ so

$$d(w, z) \leq 0.$$

The reverse inequality is also true, since d is a metric, so $d(w, z) = 0$ and hence

$$w = z.$$

What we've shown is that if E is a complete metric space, X is a closed subset of E,

$$d(f(x), f(y)) \leq q d(x, y).$$

for some $q \in (0, 1)$ and all $x, y \in X$, and the ball of radius

$$\rho = \frac{1}{1 - q} d(x_0, f(x_0))$$

about x_0 is contained in X then there is a unique $z \in X$ with

$$f(z) = z.$$

Further, $d(x_0, z) \leq \rho$. This statement, in the special case $X = E$, and with the last sentence removed, is known as the *Banach Fixed Point Theorem.*

As a quick application of this this theorem, consider the function $f\colon \mathcal{L}(E) \to \mathcal{L}(E)$ defined by

$$f(x) = I + Kx$$

where E is a Banach space, $K \in \mathcal{L}(E)$ and $\|K\| < 1$. Then

$$\begin{aligned} \|f(x) - f(y)\| &= \|I + Kx - I - Ky\| \\ &= \|K(x - y)\| \leq \|K\| \|x - y\|, \end{aligned}$$

or, in terms of the associated metric,

$$d(f(x), f(y)) \leq q d(x, y)$$

with $q = \|K\| < 1$. So the hypotheses of the Banach Fixed Point Theorem are satisfied and there is a $z \in \mathcal{L}(E)$ such that $f(z) = z$, i.e.

$$I + Kz = z$$

or

$$I = (I - K) z.$$

Similarly, using $f(x) = I + xK$, we see that there is a w in $\mathcal{L}(E)$ such that

$$I = w(I - K).$$

Then

$$w = wI = w(I - K)z = Iz = z.$$

So z is an inverse to $I - K$, which therefore must be invertible. This gives an alternate proof to the result from the section on geometric series. It's only superficially a different proof though. If you apply the proof of the Banach Fixed Point theorem to this special case then you get exactly the proof of the invertibility of $I - K$ that was presented in the earlier section.

2.10 Zorn's Lemma

Suppose S is a partially ordered set, i.e. a set with a binary relation, which we'll denote $\leq$, which satisfies the following conditions:

1. For all $x \in S$, $x \leq x$,
2. For all $x, y \in S$, if $x \leq y$ and $y \leq x$ then $x = y$, and
3. For all $x, y, z \in S$, if $x \leq y$ and $y \leq z$ then $x \leq z$.

The second property says that if $x \neq y$ then at most one of $x \leq y$ or $y \leq x$ is true, but it leaves open the possibility that neither is true. If that possibility is excluded, in other words if for every $x, y \in S$ at least one of $x \leq y$ or $y \leq x$ is true then we say that S is totally ordered.

The set $\mathbf{R}$ of real numbers, ordered by the usual less than or equal relation, is a totally ordered set. The set of subsets of $\mathbf{R}$, ordered by inclusion of sets, is partially ordered, but the set of subsets of the reals is not totally ordered, because there are pairs of sets of real numbers neither of which includes the other.

Any subset of a partially ordered set is a partially ordered set, with the relation inherited from the larger set and any subset of a totally ordered set is a totally ordered set.

If P is a subset of a partially ordered set S then a bound for P is a $y \in S$ such that for all $x \in P$ we have $x \leq y$. A greatest element of P is a $y \in P$ such that for all $x \in P$ we have $x \leq y$. A maximal element is a $y \in P$ such that $y \leq z$ for $z \in P$ only when $y = z$. Only the first of these notions contains any reference to the set S.

The following statements are easy consequences of the definitions:

- Any greatest element is maximal and a bound.
- In a totally ordered set any maximal element is a greatest element.
- There is at most one greatest element.
- If there is a greatest element then it is the only maximal element.
- Any bound which is an element of the subset is a greatest element.
- Any finite subset of a partially ordered set has a maximal element.
- Any finite subset of a totally ordered set has a greatest element.

The subset $[0, 1]$ of $\mathbf{R}$ has a greatest element, which is just 1. While this is the only greatest element and only maximal element, it is not the only bound. It is, however, the only bound lying within $[0, 1]$. The subset $(0, 1)$ of $\mathbf{R}$ has neither a greatest element nor a maximal element, but it has many bounds. $\mathbf{R}$, considered as a subset of itself has no bound, no maximal element and no greatest element.

The set of subsets of $\mathbf{R}$, with the inclusion relation, has a greatest element, namely $\mathbf{R}$ itself, since every subset is contained in $\mathbf{R}$. This greatest element is both a bound and a maximal element. The set of proper subsets of $\mathbf{R}$ has maximal elements but no greatest element. The maximal elements are all of the form $\mathbf{R} - \{x\}$ for some $x \in \mathbf{R}$. The set of proper subsets, considered as a subset of the set of all subsets, has a bound, namely $\mathbf{R}$.

Zorn's Lemma is the statement that if every totally ordered subset of a partially ordered subset has a bound then the partially ordered set has a maximal element. Zorn's Lemma follows from the usual axioms of set theory. In fact it is equivalent to one of them, the Axiom of Choice, in the sense that if we replace that axiom with Zorn's Lemma then we can prove the Axiom of Choice as a theorem. If, instead, we simply drop the Axiom of Choice then we can prove neither of them.

The main point of Zorn's lemma is to "construct" maximal objects of various types. For example, given any linearly independent set A in a vector space V we can consider the set P of all linearly independent sets L of vectors with $A \subseteq L$, ordered by inclusion of sets. Let T be a totally ordered subset of P and let $M = \cup_{L \in T} L$. Then M is a linearly independent subset of V. To see this, observe that if

$$\sum_{j=1}^{n} \alpha_j x_j = 0$$

and $a = (\alpha_1, \ldots, \alpha_n) \in \mathbf{K}^n$, $x_1, \ldots, x_n \in M$ then $x_j \in L_j$ for some $L_j \in T$. Since T is totally ordered we can, after possibly rearranging the order in the finite sum above, assume that $L_1 \subseteq L_2 \subseteq \ldots \subseteq L_n$. But then $x_1, \ldots, x_n \in L_n$. Since L_n is a linearly independent subset it follows that $a = 0$. So any linear relation in M is trivial. In other words M is a linearly independent set. It contains A, so $M \in P$. It also contains each $L \in T$, so its a bound for T. So the hypothesis of Zorn's Lemma, that every totally ordered subset is bounded, is satisfied. There therefore is a maximal element $B \in P$. This B must span V, because there would otherwise be an $x \in V$ which is not a linear combination of elements of B and the set $B \cup \{x\}$ would be an element of P larger than the maximal element B. So B is a spanning linearly independent set, i.e. a basis. This shows that any linearly independent set can be extended to a basis. One can similarly use Zorn's Lemma to show that every spanning set has a subset which is a basis.

The word 'basis' is used here in the sense which is standard in Linear Algebra. To avoid confusion with other notions of basis it is standard to call such bases *Hamel bases* in Functional Analysis. Hamel bases are, as it turns out, almost completely useless in Functional Analysis.

2.11 The Inverse Function Theorem

Suppose U is an open subset of a normed space E and f is a function from U to a normed space F. A bounded linear transformation A is said to be the derivative of f at x if

$$\lim_{y \to x} \frac{\|f(y) - f(x) - A(y-x)\|}{\|y-x\|} = 0.$$

If E is not zero dimensional, as we shall henceforth assume, then there is at most one derivative. The derivative of f at x is then denoted $(Df)(x)$. If there is a derivative at x then we say that f is differentiable at x. If there is a derivative all $x \in U$ then we say simply that f is differentiable. There is then a function which associates to each $x \in U$ the derivative of f at x, which is an element of $\mathcal{L}(E,F)$. This function is denoted by Df, which is consistent with the notation $(Df)(x)$ adopted above. f is said to be continuously differentiable if $Df\colon U \to \mathcal{L}(E,F)$ is a continuous function.

Suppose now that Df is bounded on a convex set U, i.e. that

$$\|(Df)(y)\| \le \lambda$$

for all $y \in U$. We will now show that

$$\|f(z) - f(x)\| \le \lambda \|z - x\|$$

for $x, z \in U$.

Set

$$g(t) = f((tz + (1-t)x))$$

for $0 \le t \le 1$ and

$$h(s,t) = (t-s)^{-1}(g(t) - g(s))$$

for $0 \le s < t \le 1$. Then, for $0 \le r < s < t \le 1$ we have

$$h(r,t) = \frac{s-r}{t-r} h(r,s) + \frac{t-s}{t-r} h(s,t)$$

and hence

$$\|h(r,t)\| \le \left|\frac{s-r}{t-r}\right| \|h(r,s)\| + \left|\frac{t-s}{t-r}\right| \|h(s,t)\|$$

and

$$\|h(r,t)\| \le \mu \max\left(\|h(r,s)\|, \|h(s,t)\|\right)$$

where

$$\mu = \left(\left|\frac{s-r}{t-r}\right| + \left|\frac{t-s}{t-r}\right|\right) = 1.$$

So at least one of $h(r,s)$ or $h(s,t)$ has a larger norm than $h(r,t)$. We then define a sequence of intervals $[a_n, b_n]$ inductively as follows. $[a_1, b_1]$ is $[0,1]$ and $[a_{k+1}, b_{k+1}]$ is either $[a_k, (a_k + b_k)/2]$ or $[(a_k + b_k)/2, b_k]$, chosen in such a way that

$$\|h(a_k, b_k)\| \le \|h(a_{k+1}, b_{k+1})\|.$$

The sequence $(a_1, a_2, \ldots)$ is bounded and monotone increasing. The sequence $(b_1, b_2, \ldots)$ is bounded and monotone decreasing. Both therefore have limits. Since $b_n - a_n = 2^{1-n}$ these limits are the same. Call their common limit q and let

$$y = qz + (1-q)x.$$

For any $\epsilon > 0$ there is then, by the definition of Df, a $\delta > 0$ such that $\|w - y\| < \delta$ implies

$$\frac{\|f(w) - f(y) - (Df)(y)(w-y)\|}{\|w-y\|} < \epsilon,$$

from which it follows that

$$\frac{\|f(w) - f(y)\|}{\|w-y\|} \le \left(\|(Df)(y)\| + \epsilon\right).$$

For n sufficiently large, both

$$w = a_n z + (1 - a_n)x$$

and

$$w = b_n z + (1 - b_n)x$$

are within a distance δ of y. In fact

$$n > 1 + \log\left(\frac{\delta}{\|z - x\|}\right) / \log 2$$

suffices. For such n then

$$\|h(a_n, q)\| \leq (\|(Df)(y)\| + \epsilon) \|z - x\|$$

and

$$\|h(q, b_n)\| \leq (\|(Df)(y)\| + \epsilon) \|z - x\|.$$

For any such n we apply the inequality derived above with $r = a_n$, $s = q$ and $t = b_n$,

$$\|h(a_n, b_n)\| \leq \max(\|h(a_n, q)\|, \|h(q, b_n)\|)$$

and hence

$$\|h(a_n, b_n)\| \leq (\|(Df)(y)\| + \epsilon) \|z - x\|$$

It then follows that

$$\|f(z) - f(x)\| = \|h(a_1, b_1)\| \leq (\lambda + \epsilon) \|z - x\|.$$

Since this holds for all $\epsilon > 0$ we conclude that

$$\|f(z) - f(x)\| \leq \lambda \|z - x\|.$$

As a consequence of the inequality above, we note that if X is an open convex subset of a Banach space E, $f: X \to E$ is continuously differentiable and

$$\|(Df)(y)\| \leq q < 1$$

for $y \in X$ then f is a contraction mapping on X.

The local version of the Inverse Function Theorem states that if E, F are Banach spaces, $U \subseteq E$ is open $f: U \to F$ is continuously differentiable, $x_0 \in U$, $f(x_0) = y_0$ and $(Df)(x_0)$ is invertible in $\mathcal{L}(E, F)$ then there are open sets $V \subseteq E$ and $W \subseteq F$ and a continuously differentiable $g: W \to V$ such that

1. $V \subseteq U$,
2. $x_0 \in V$,
3. $y_0 \in W$,
4. $f(V) = W$
5. $g(f(x)) = x$ for all $x \in V$,
6. $f(g(y)) = y$ for all $y \in W$,
7. $(Dg)(y) = (Df)(g(y))^{-1}$.

First note that $(Df)(x_0)^{-1}(Df)(x)$ is continuous and equal to I when $x = x_0$ so there is a $\delta > 0$ such that $\|x - x_0\| < \delta$ implies $x \in U$ and

$$\left\|(Df)(x_0)^{-1}(Df)(x) - I\right\| < \frac{1}{2}.$$

Let X be the closed ball of radius $\delta/2$ about x_0,

$$X = \left\{x \in E: \|x - x_0\| \leq \frac{\delta}{2}\right\}.$$

So X is a closed convex set on which

$$\left\|(Df)(x_0)^{-1}(Df)(x) - I\right\| < \frac{1}{2}.$$

Define

$$k(x) = x + (Df)(x_0)^{-1}(y - f(x))$$

for $x \in U$. This k depends, of course, on a choice of $y \in F$. For the moment we'll leave this y arbitrary. Then k is continuously differentiable in U. In fact

$$(Dk)(x) = I - (Df)(x_0)^{-1}(Df)(x).$$

With X as above,

$$\|(Dk)(x)\| \leq \frac{1}{2}$$

for $x \in X$. k is then a contraction mapping in X. By the version of the Banach Fixed Point Theorem proved earlier we then have a unique fixed point in X provided that the ball of radius

$$\rho = \frac{x_0 - k(x_0)}{1 - 1/2} = 2\left\|(Df)(x_0)^{-1}(y - y_0)\right\|$$

is contained in X, i.e. provided that

$$\left\|(Df)(x_0)^{-1}(y-y_0)\right\| \le \frac{\delta}{4}.$$

Looking at the definition of k, we see that x is a fixed point of k if and only if

$$y = f(x).$$

We define

$$W = \left\{y \in F \colon \left\|(Df)(x_0)^{-1}(y-y_0)\right\| < \frac{\delta}{4}\right\}$$

and define $g(y)$ for $y \in W$ to be the unique fixed point x of k whose existence we've just shown. We define $V = g(W)$.

By construction, $f(g(y)) = y$ for $y \in W$. Also, if $x \in V$ then $x = g(y)$ for some $y \in W$. Then $f(x) = f(g(y)) = y$ and $g(f(x)) = g(y) = x$. It remains to show that g is continuously differentiable.

Let

$$h(x) = (Df)(x_0)^{-1}f(x) - x.$$

This is a continuously differentiable function on U with derivative

$$Dh = (Df)(x_0)^{-1}Df - I$$

and therefore have

$$\|(Dh)(x)\| \le \frac{1}{2}$$

for x in the closed convex set X and then

$$\|h(x_1) - h(x_2)\| \le \frac{1}{2}\|x_1 - x_2\|$$

for x_1, x_2 in X. Now

$$x_1 - x_2 = (Df)(x_0)^{-1}(f(x_1) - f(x_2)) - h(x_1) - h(x_2)$$

and hence

$$\begin{aligned}\|x_1 - x_2\| &\le \|(Df)(x_0)^{-1}(f(x_1) - f(x_2))\| \\ &\quad + \|h(x_1) - h(x_2)\|, \\ &\le \|(Df)(x_0)^{-1}(f(x_1) - f(x_2))\| \\ &\quad + \frac{1}{2}\|x_1 - x_2\|,\end{aligned}$$

and

$$\|x_1 - x_2\| \le 2\|(Df)(x_0)^{-1}(f(x_1) - f(x_2))\|.$$

Applying this to

$$x_1 = g(y_1)$$

and

$$x_2 = g(y_2)$$

for $y_1, y_2 \in W$ gives

$$\begin{aligned}\|g(y_1) - g(y_2)\| &\le 2\|(Df)(x_0)^{-1}(y_1 - y_2)\| \\ &\le 2\|(Df)(x_0)^{-1}\|\|(y_1 - y_2)\|.\end{aligned}$$

This shows that g is continuous. A slightly more sophisticated version of the same argument shows that it is in fact differentiable, with

$$(Dg)(y) = (Df)(g(y))^{-1}.$$

This is the composition of g, Df and inversion, all of which are known to be continuous, so Dg is continuous as well.

It would be nice to have a global version of the Inverse Function Theorem, showing that any continuously differentiable function with an invertible derivative has an inverse defined on its entire range. Unfortunately the functions $\exp\colon \mathbf{C} \to \mathbf{C}$ shows that this is impossible. The best substitute we can find is a maximal version, proved using the local version and Zorn's Lemma. If E, F are Banach spaces, $U \subseteq E$ is open $f\colon U \to F$ is continuously differentiable, $x_0 \in U$, $f(x_0) = y_0$ and $(Df)(x_0)$ is invertible in $\mathcal{L}(E, F)$ then there are open sets $V \subseteq U$ and $W \subseteq F$ and a continuously differentiable $g\colon W \to V$ such that

1. $x_0 \in V$,
2. $y_0 \in W$,
3. $f(V) = W$
4. $g(f(x)) = x$ for all $x \in V$,
5. $f(g(y)) = y$ for all $y \in W$,
6. $(Dg)(y) = (Df)(g(y))^{-1}$, and
7. if $\tilde{V} \subseteq E$, $\tilde{W} \subseteq F$ are open sets and $\tilde{g}\colon \tilde{W} \to \tilde{V}$ a continuously differentiable function with

(a) $x_0 \in \tilde{V}$,

(b) $y_0 \in \tilde{W}$,

(c) $f(\tilde{V}) = \tilde{W}$

(d) $\tilde{g}(f(x)) = x$ for all $x \in \tilde{V}$,

(e) $f(\tilde{g}(y)) = y$ for all $y \in \tilde{W}$,

then W is not a proper subset of $\tilde{W}$.

2.12 The Spaces $L^p(I)$

In this section we suppose that I is an interval in $\mathbf{R}$, whether finite, semi-infinite or infinite. In fact for most of the results of this section we can take I to be any measurable subset of $\mathbf{R}^n$ or even something more general than that, but for purposes of illustration intervals are enough. All integrals are Lebesgue integrals.

For $1 \leq p < \infty$ we define $F^p(I)$ to be the set of measurable functions f on I for which $|f|^p$ is integrable on I. For such functions

$$q(f) = \int_I |f|^p$$

exists. In fact it defines a semi-norm. All the properties of a semi-norm are obvious except for the triangle inequality, which follows from the integral version of the Minkowski Inequality.

The semi-norm q is not, unfortunately, a norm. $q(f) = 0$ if and only if the set of points x where $f(x) \neq 0$ has measure 0. Let N be the space of such functions. As saw in Question 1 of Assignment 2, and as we can easily check directly, N is a vector subspace. We can then form the quotient spaces $L^p(I) = F^p(I)/N$. As we also saw in the same assignment question, q then defines a norm on the quotient space $L^p(I)$. This norm will be written as $\|f\|$ or, when we need to distinguish different values of p, as $\|f\|_p$.

The disadvantage of working with the quotient space is that it doesn't make sense to evaluate an element of $L^p(I)$ at a point, since points have measure zero and thus two elements of an N-coset may take different values at a point. This is more than compensated for though by the advantages of working with a norm rather than a semi-norm. To make things less awkward we often refer to functions f as elements of $L^p(I)$ rather than as coset representatives of cosets of N. As usual in measure theory, we refer to statements which are valid except possibly on a set of measure zero as holding 'almost everywhere.' So functions f and g represent the same element of $L^p(I)$ if and only if $f = g$ almost everywhere.

A useful fact is that absolute convergence in $L^p(I)$ implies pointwise convergence almost everywhere. More precisely, if $(f_1, f_2, \ldots)$ is a sequence in $L^p(I)$ such that

$$\sum_{j=1}^{\infty} \|f_j\| < \infty$$

then there is a $g \in L^p(I)$ such that

$$\sum_{j=1}^{\infty} f_j = g$$

in $L^p(I)$ and

$$\sum_{j=1}^{\infty} f_j(x) = g(x)$$

almost everywhere in I.

To see this, define

$$h_n(x) = \sum_{j=1}^{n} |f_j(x)|, \quad k(x) = \lim_{n \to \infty} h_n(x).$$

The sequence $(h_1(x), h_2(x), \ldots)$ is monotone increasing and converges to $k(x)$. From this it follows that the sequence $(h_1(x)^p, h_2(x)^p, \ldots)$ is monotone increasing and converges to $k(x)^p$. It follows from the Monotone Convergence Theorem that

$$\lim_{n \to \infty} \int_I h_n^p = \int_I k^p.$$

By Minkowski's Inequality

$$\|h_n\| \leq \sum_{j=1}^{n} \|f_j\| \leq \sum_{j=1}^{\infty} \|f_j\| < \infty.$$

It follows that $k(x)$ is finite for almost all x, i.e. that the sum of $f_j(x)$ is absolutely convergent for almost all x. It therefore makes sense to define

$$g(x) = \sum_{j=1}^{\infty} f_j(x).$$

Also $|g(x)| \leq k(x)$. The fact that this converges only almost everywhere is irrelevant, because we are defining an element of $L^p(I)$, not of $F^p(I)$. Then

$$\left| g(x) - \sum_{j=1}^{n} f_j(x) \right|^p \leq 2^p |k(x)|^p.$$

Since we already know that k^p is integrable we can apply the Lebesgue Dominated Convergence Theorem to say that

$$\lim_{n\to\infty} \int_I \left| g - \sum_{j=1}^{n} f_j \right|^p = \int_I \lim_{n\to\infty} \left| g - \sum_{j=1}^{n} f_j \right|^p = 0.$$

In other words

$$\lim_{n\to\infty} \|g - \sum_{j=1}^{n} f_j\|^p = 0.$$

Hence

$$\lim_{n\to\infty} \|g - \sum_{j=1}^{n} f_j\| = 0$$

and

$$\lim_{n\to\infty} \sum_{j=1}^{n} f_j = g$$

in $L^p(I)$.

We can use this fact to show that $L^p(I)$ is complete, i.e. that it is a Banach space. Suppose that $(f_1, f_2, \ldots)$ is a Cauchy sequence in $L^p(I)$. In other words, suppose there is for each $\epsilon > 0$ an N such that

$$\|f_j - f_k\| < \epsilon$$

whenever $j, k > N$. This holds in particular for

$$\epsilon = \frac{1}{2^m},$$

so there is an N_m such that for $j, k > N_m$ we have

$$\|f_j - f_k\| \leq \frac{1}{2^m}.$$

Let

$$g_m = f_{l_m}$$

where

$$l_m = \max(m, N_{m-1}, N_m) + 1$$

and

$$h_m = g_{m+1} - g_m.$$

Then

$$\|h_m\| = \|g_{m+1} - g_m\| < \frac{1}{2^m}.$$

So

$$\sum_{m=1}^{\infty} \|h_m\| < \infty.$$

But as we've seen this implies the convergence of

$$\lim_{n\to\infty} \sum_{m=1}^{n} h_m$$

in $L^p(I)$. But

$$g_n = g_1 + \sum_{m=1}^{n-1} h_m$$

so $(g_1, g_2, \ldots) = (f_{l_1}, f_{l_2}, \ldots)$ converges in $L^p(I)$. A Cauchy sequence with a convergent subsequence is convergent, so $(f_1, f_2, \ldots)$ must be convergent.

3 Inner Product and Hilbert Spaces

3.1 Inner Product Spaces

An inner product on a vector space E is a function $q\colon E \times E \to \mathbf{K}$ such that for all $\alpha \in \mathbf{K}$ and $x, y, z \in E$ we have

1. $q(x + y, z) = q(x, z) + q(y, z)$,
2. $q(\alpha x, y) = \alpha q(x, y)$,
3. $q(x, y) = \overline{q(y, x)}$, and
4. $q(x, x) \geq 0$
5. $q(x, x) = 0$ only if $x = 0$.

The bar in the third property above denotes complex conjugation. If $\mathbf{K} = \mathbf{R}$ then it has no effect. The second to last of these properties requires some comment, since in general $q(x, y) \in \mathbf{K}$ and $\mathbf{K}$ is allowed to be $\mathbf{C}$, which doesn't have a $\geq$ relation. But it follows from the third property that

$$q(x, x) = \overline{q(x, x)}$$

so $q(x, x) \in \mathbf{R}$ even if $\mathbf{K} = \mathbf{C}$.

A vector space together with an inner product is called an inner product space. Unless we are considering multiple inner products on the same space we will write $(x|y)$ in place of $q(x, y)$, so the properties above become

1. $(x + y|z) = (x|z) + (y|z)$,
2. $(\alpha x|y) = \alpha\,(x|y)$,
3. $(x|y) = \overline{(y|x)}$, and
4. $(x|x) \geq 0$
5. $(x|x) = 0$ only if $x = 0$.

Some useful additional properties which follow immediately from these are

$$(x|y + z) = (x|y) + (x|z)$$

and

$$(x|\alpha y) = \overline{\alpha}\,(x|y)\,.$$

Given an inner product, we can define a norm by

$$\|x\| = \sqrt{(x|x)}.$$

The first two properties of norms are clearly satisfied by this, but the triangle inequality requires more work. We start by noting that if

$$z = \alpha x + \beta y$$

where $\alpha, \beta \in \mathbf{K}$ and $x, y \in E$ then we have

$$\begin{aligned}
(z|z) &= (\alpha x + \beta y|\alpha x + \beta y)\\
&= (\alpha x|\alpha x + \beta y) + (\beta y|\alpha x + \beta y)\\
&= \alpha\,(x|\alpha x + \beta y) + \beta\,(y|\alpha x + \beta y)\\
&= \alpha\overline{(\alpha x + \beta y|x)} + \beta\overline{(\alpha x + \beta y|y)}\\
&= \alpha\overline{[(\alpha x|x) + (\beta y|x)]} + \beta\overline{[\alpha\,(x|y) + \beta\,(y|y)]}\\
&= \alpha\overline{\alpha}\overline{(x|x)} + \alpha\overline{\beta}\overline{(y|x)} + \beta\overline{\alpha}\overline{(x|y)} + \beta\overline{\beta}\overline{(y|y)}\\
&= \alpha\overline{\alpha}\,(x|x) + \alpha\overline{\beta}\,(x|y) + \beta\overline{\alpha}\,(y|x) + \beta\overline{\beta}\,(y|y)\,.
\end{aligned}$$

If

$$\alpha = (y|x)\,, \quad \beta = -\,(x|x)$$

then

$$\begin{aligned}
(z|z) &= (y|x)\,\overline{(y|x)}\,(x|x) - (y|x)\,\overline{(x|x)}\,(x|y)\\
&\quad - (x|x)\,\overline{(y|x)}\,(y|x) + (x|x)\,\overline{(x|x)}\,(y|y)\\
&= (y|x)\,(x|y)\,(x|x) - (y|x)\,(x|x)\,(x|y)\\
&\quad - (x|x)\,(x|y)\,(y|x) + (x|x)\,(x|x)\,(y|y)\\
&= (x|x)\,[(x|x)\,(y|y) - (x|y)\,(y|x)].
\end{aligned}$$

Similarly, if

$$\alpha = (y|y)\,, \quad \beta = -\,(x|y)$$

then

$$(z|z) = (y|y)\,[(x|x)\,(y|y) - (x|y)\,(y|x)].$$

We know that $(x|x) \geq 0$ and $(y|y) \geq 0$ for all $x, y \in E$. If either $(x|x) > 0$ or $(y|y) > 0$ then it follows that

$$(x|x)\,(y|y) - (x|y)\,(y|x) \geq 0.$$

If both $(x|x) = 0$ and $(y|y) = 0$ then we take

$$\alpha = 1, \quad \beta = -\,(x|y)\,.$$

This gives

$$(z|z) = -2\,(x|y)\,(y|x) = -2|\,(x|y)\,|^2.$$

From the non-negativity of $(z|z)$ it then follows that $|\,(x|y)\,|^2 \leq 0$ and then that

$$|\,(x|y)\,|^2 = 0,$$

and hence

$$|\,(x|y)\,|^2 \leq (x|x)\,(y|y)\,.$$

Thus

$$|\,(x|y)\,|^2 \leq (x|x)\,(y|y)\,.$$

for all $x, y \in E$. For future reference we note that we never used the last property of the inner product in this proof.

Taking square roots,

$$|\,(x|y)\,| \leq \sqrt{(x|x)\,(y|y)}$$

or

$$|\,(x|y)\,| \leq \sqrt{\|x\|\|y\|}.$$

This is known as the Cauchy-Schwarz inequality. The various other Cauchy-Schwarz inequalities we've seen are all special cases of this one.

Using the Cauchy-Schwarz inequality,

$$\begin{aligned}
\|x+y\|^2 &= (x+y|x+y) \\
&= (x|x) + (x|y) + (y|x) + (y|y) \\
&= (x|x) + (x|y) + \overline{(x|y)} + (y|y) \\
&= \|x\|^2 + 2\operatorname{Re}(x|y) + \|y\|^2 \\
&\le \|x\|^2 + 2\|x\|\|y\| + \|y\|^2 \\
&\le (\|x\| + \|y\|)^2 .
\end{aligned}$$

Taking square roots,

$$\|x+y\| \le \|x\| + \|y\|.$$

That's the triangle inequality, so this completes the proof that what we've called the norm is in fact a norm.

The inner product is continuous in the sense that if $x_n \to w$ and $y_n \to z$ then $(x_n|y_n) \to (w|z)$. To see this, note that

$$\begin{aligned}
|\,(x_n|y_n) - (w|z)\,| &= |\,(x_n - w|z) + (x_n|y_n - z)\,| \\
&\le \|x_n - w\|\|z\| + \|x_n\|\|y_n - z\|.
\end{aligned}$$

Since $(x_1, x_2, \ldots)$ is a convergent sequence it is bounded, i.e.

$$\|x_n\| \le \mu$$

for some $\mu \in \mathbf{R}^+$. Choose a $\nu \in \mathbf{R}^+$ such that $\|z\| \le \nu$. $\nu = \|z\|$ will do if $z \ne 0$ and any $\nu \in \mathbf{R}^+$ will do if $z = 0$. For $\epsilon > 0$ the quantities

$$\frac{\epsilon}{2\mu}, \quad \frac{\epsilon}{2\nu}$$

are both positive so there are $N_1, N_2 \in \mathbf{Z}^+$ such that $j > N_1$ implies

$$\|x_n - w\| < \frac{\epsilon}{2\nu}$$

and $j > N_2$ implies

$$\|y_n - z\| < \frac{\epsilon}{2\mu}.$$

Taking $N = \max(N_1, N_2)$, if $j > N$ then

$$|\,(x_n|y_n) - (w|z)\,| < \epsilon.$$

This shows that $(x_n|y_n) \to (w|z)$.

Examples of inner product spaces include $l^2(n)$ with the inner product

$$(x|y) = \sum_{j=1}^{n} \xi_j \overline{\eta_j},$$

l^2 with the inner product

$$(x|y) = \sum_{j=1}^{\infty} \xi_j \overline{\eta_j},$$

and $L^2(I)$ with the inner product

$$\int_I x\overline{y}.$$

In each case the norm induced by the inner product is the usual norm. Another example, generalising $l^2(n)$ and l^2, is the set $l^2(T)$ where T is an arbitrary set and $l^2(T)$ consists of functions $x\colon T \to \mathbf{K}$ such that $x(t) \ne 0$ for at most countably many $t \in T$ and

$$\sum_{t\in T} |x(t)|^2 < \infty.$$

The inner product is

$$(x|y) = \sum_{t\in T} x(t)\overline{y(t)}.$$

An inner product space which is complete as a metric space or, equivalently, is a Banach space as a normed space is called a Hilbert space. All of the examples of inner product spaces above are Hilbert spaces. To get an example of an inner product space which is not a Hilbert space we can take $C([a,b])$ with the inner product it inherits as a subspace of $L^2([a,b])$.

3.2 Orthogonality

We say that vectors x and y in an inner product space are orthogonal if

$$(x|y) = 0.$$

There are two useful ways to generalise this. We say that subsets M and N are orthogonal if $(x|y) = 0$

for all $x \in M$ and $y \in N$. We also say that a set S of vectors are pairwise orthogonal if $(x|y) = 0$ for all distinct $x, y \in S$. We'll refer to such a subset as an orthogonal subset. If $S = \{x_1, \ldots x_n\}$ is such a subset then the distributive law shows that

$$\left\| \sum_{j=1}^{n} x_j \right\|^2 = \sum_{j=1}^{n} \|x_j\|^2.$$

There are no cross terms because of the orthogonality assumption. We can generalise this is countable subsets $S = \{x_1, x_2, \ldots\}$, with

$$\left\| \sum_{j=1}^{\infty} x_j \right\|^2 = \sum_{j=1}^{\infty} \|x_j\|^2.$$

More precisely, if we assume that $\sum_{j=1}^{\infty} x_j$ converges then so does $\sum_{j=1}^{\infty} \|x_j\|^2$ and the equation holds. Let

$$s_n = \sum_{j=1}^{n} x_j$$

and

$$z = \lim_{n \to \infty} s_n$$

which, by assumption, exists. Then the finite version shows that

$$\sum_{j=1}^{n} \|x_j\|^2 = \|s_n\|^2 = (s_n|s_n)$$

The right hand side tends to $(z|z)$ by the continuity of the inner product, so the left hand side tends to that as well.

An orthogonal subset is called orthonormal if all its elements have norm 1. If S is a maximal orthonormal set and $\{y\}$ is orthogonal to S then $y = 0$, since otherwise $S \cup \{y/\|y\|\}$ would be a strictly larger orthonormal set. Conversely, if 0 is the only vector orthogonal to S then S is maximal. An examples of a maximal orthonormal subsets is

$$\{e_1, e_2, \ldots, e_n\} \subseteq l^2(n),$$

where e_j is all zeroes except in the j'th position, which is a 1. Similarly

$$\{e_1, e_2, \ldots\} \subseteq l^2$$

is a maximal orthonormal set. Maximality is clear in both cases because if $x = (\xi_1, \xi_2, \ldots)$

$$(x|e_j) = \xi_j$$

so if x is orthogonal to all the e's then it must be the zero sequence. The set

$$\left\{ \frac{1}{\sqrt{2}}, \cos(\pi t), \sin(\pi t), \cos(2\pi t), \sin(2\pi t), \ldots \right\}$$

is a maximal orthonormal set in $L^2([-1,1])$. Orthonormality is an easy calculation but maximality is less clear in this case. Another maximal orthonormal set for $L^2([-1,1])$ is $p_1, p_2, \ldots$

$$p_j = \sqrt{j - \frac{1}{2}} P_{j-1}$$

and P_n is the n'th Legendre polynomial.

Suppose $S = \{u_1, \ldots u_n\}$ is an orthonormal set in an inner product space E, but not necessarily a maximal one. Suppose that $x \in E$ and $\alpha_1, \ldots, \alpha_n \in \mathbf{K}$. Let

$$z = x - \sum_{j=1}^{n} \alpha_j u_j.$$

Then

$$\begin{aligned}
\|z\|^2 &= (z|z) \\
&= (x|x) - \sum_{j=1}^{n} \alpha_j \,(u_j|x) \\
&\quad - \sum_{k=1}^{n} \overline{\alpha}_k \,(x|u_k) + \sum_{j=1}^{n} \sum_{k=1}^{n} \alpha_j \overline{\alpha}_k \,(u_j|u_k) \\
&= (x|x) - \sum_{j=1}^{n} \alpha_j \overline{(x|u_j)} \\
&\quad - \sum_{j=1}^{n} \overline{\alpha}_j \,(x|u_j) + \sum_{j=1}^{n} \alpha_j \overline{\alpha}_j \\
&= (x|x) + \sum_{j=1}^{n} [\alpha_j - (x|u_j)][\overline{\alpha_j} - \overline{(x|u_j)}] \\
&\quad - \sum_{j=1}^{n} (x|u_j) \,\overline{(x|u_j)} \\
&= \|x\|^2 - \sum_{j=1}^{n} |\,(x|u_j)\,|^2 + \sum_{j=1}^{n} |\alpha_j - (x|u_j)\,|^2.
\end{aligned}$$

It follows that

$$\|z\|^2 \geq \|x\|^2 - \sum_{j=1}^{n} |\,(x|u_j)\,|^2$$

with equality if and only if

$$\alpha_j = (x|u_j)$$

for each j. Taking $\alpha_j = (x|u_j)$ in the equation above gives

$$\left\| x - \sum_{j=1}^{n} (x|u_j)\, u_j \right\|^2 = \|x\|^2 - \sum_{j=1}^{n} |\,(x|u_j)\,|^2,$$

which is called Bessel's identity. Since the left hand side is non-negative this implies the Bessel inequality

$$\sum_{j=1}^{n} |\,(x|u_j)\,|^2 \leq \|x\|^2.$$

3.3 Orthogonalisation

If $\{x_1, x_2, \ldots\}$ is a linearly independent subset of an inner product space E then we can find an orthonormal set $\{u_1, u_2, \ldots\}$ as follows. Set

$$y_k = x_k - \sum_{j=1}^{k-1} (x_k|u_j)\, u_j.$$

and

$$u_k = \frac{y_k}{\|y_k\|}.$$

The sum is empty if $k = 1$. By induction, starting with $k = 1$, we have the following facts

1. The vectors $u_1, \ldots, u_k$ are linear combinations of the vectors $x_1, \ldots, x_k$ and vice versa. In particular they have the same span.

2. If $j < k$ then u_j is orthogonal to y_k and hence to u_k.

3. $y_k \neq 0$, so u_k exists and has norm 1.

These are all straightforward. Suppose they're all true for a given value of k. Then y_{k+1} is non-zero since the equation defining it would otherwise give x_{k+1} as a linear combination of $u_1, \ldots, u_k$ and hence of $x_1, \ldots, x_k$, contradicting our linear independence assumption. u_{k+1} is given explicitly as a linear combination of $u_1, \ldots, u_k, x_{k+1}$ and by the inductive assumption $u_1, \ldots, u_k$ are linear combinations of $x_1, \ldots, x_k$, so u_{k+1} is a linear combination of $x_1, \ldots, x_k, x_{k+1}$. Also

$$x_{k+1} = \|y_{k+1}\| u_{k+1} + \sum_{j=1}^{k} (x_{k+1}|u_j)\, u_j,$$

so x_{k+1} is a linear combination of $u_1, \ldots u_l, u_{k+1}$. Taking inner products with u_l in the equation

$$y_{k+1} = x_{k+1} - \sum_{j=1}^{k} (x_{k+1}|u_j)\, u_j$$

gives

$$(y_{k+1}|u_l) = (x_{k+1}|u_l) - \sum_{j=1}^{k} (x_{k+1}|u_j)\,(u_j|u_l)$$

for $l < k + 1$. In the sum on the right

$$(u_j|u_l) = \begin{cases} 1 & \text{if } j = l \\ 0 & \text{if } j \neq l \end{cases},$$

by the inductive hypotheses, so

$$(y_{k+1}|u_l) = (x_{k+1}|u_l) - (x_{k+1}|u_l) = 0.$$

The algorithm which finds u_1, u_2, …from x_1, x_2, …is known as the Gram-Schmidt procedure. It can also be applied to check whether the set $\{x_1, x_2, \ldots\}$ is linearly independent. Simply run the algorithm and if $y_k = 0$ at some stage then $\{x_1, \ldots, x_{k-1}\}$ are linearly independent but $\{x_1, \ldots, x_k\}$ are not. An orthonormal set is automatically linearly independent because

$$\left\| \sum_{j=1}^{n} \alpha_j u_j \right\|^2 = \sum_{j=1}^{n} |\alpha_j|^2,$$

which is zero if and only if all the α's are.

3.4 Orthogonal Complements

Suppose E is a Hilbert space, $x \in E$ and Y is a closed convex subset of E, for example a closed subspace. Then there is a unique $y \in Y$ which minimises $\|x - y\|$. In fact this holds in the more general setting where E is an inner product space and Y is a convex subset of E which is complete as a metric space. For the uniqueness we don't even need the completeness assumption. We now prove this.

The identity

$$\|u+v\|^2 + \|u-v\|^2 = 2\|u\|^2 + 2\|v\|^2,$$

which is known as the Parallelogram Identity, holds in any inner product space, as shown in Assignment 4. Applied to

$$u = x - w, \quad v = x - z$$

this gives

$$\begin{aligned} \|w-z\|^2 &= 2\|x-w\|^2 + 2\|y-z\|^2 \\ &\quad - 4\left\| x - \frac{w+z}{2} \right\|^2 \\ &= 2\left(\|x-w\|^2 - \left\| x - \frac{w+z}{2} \right\|^2 \right) \\ &\quad + 2\left(\|x-z\|^2 - \left\| x - \frac{w+z}{2} \right\|^2 \right) \end{aligned}$$

If $w, z \in Y$ then

$$\frac{w+z}{2} \in Y$$

by convexity. If w, z are minimisers of $\|x - y\|$ then

$$\|x-w\|^2 \le \left\| x - \frac{w+z}{2} \right\|^2,$$

$$\|x-z\|^2 \le \left\| x - \frac{w+z}{2} \right\|^2,$$

and hence

$$\|w-z\|^2 \le 0$$

from which it follows that $w = z$. This establishes the uniqueness part of the statement.

Set

$$\lambda = \inf_{y \in Y} \|x-y\|^2.$$

Since $\lambda + 1/n$ is greater than λ it is not a lower bound for $\|x-y\|^2$ and hence there is then a $y_n \in Y$ such that

$$\|x - y_n\|^2 < \lambda + \frac{1}{n}.$$

The convexity of Y implies that

$$\frac{y_j + y_k}{2} \in Y$$

and hence

$$\left\| x - \frac{y_j + y_k}{2} \right\| \ge \lambda$$

We already know that

$$\|x - y_j\|^2 < \lambda + \frac{1}{j}, \quad \|x - y_j\|^2 < \lambda + \frac{1}{k}$$

and so

$$\|y_j - y_k\|^2 < \frac{2}{j} + \frac{2}{k}.$$

If $\epsilon > 0$ then we choose any $N > 4/\epsilon^2$ and we have that $j, k > N$ implies

$$\|y_j - y_k\| < \epsilon.$$

So the sequence is Cauchy and hence, by our completeness assumption, tends to an element $z \in Y$. Taking limits in the inequality

$$\|x - y_n\|^2 < \lambda + \frac{1}{n}$$

shows that

$$\|x - z\|^2 \le \lambda$$

The reverse inequality follows from the definition of λ and the fact that $z \in Y$. So

$$\|x - z\| = \inf_{y \in Y} \|x - y\|$$

and therefore z is the minimiser we're looking for.

If Y is a subspace of an inner product space E, $x \in E$ and the function $\|x - y\|$ for $y \in Y$ has a minimiser z then $x - z$ is orthogonal to Y. We show this as follows.

$$\begin{aligned} \|x-z\|^2 &\le \|x - (z + \alpha y)\|^2 \\ &= \|x-z\|^2 - \overline{\alpha}\,(x-z|y) \\ &\quad - \alpha\,(y|x-z) + \alpha\overline{\alpha}\|y\|^2 \end{aligned}$$

for all $y \in Y$ and $\alpha \in \mathbf{K}$, since $z + \alpha y \in Y$. In particular this holds for any $y \neq 0$ and

$$\alpha = \frac{(x - z|y)}{(y|y)}$$

and so

$$\|x - z\|^2 \leq \|x - z\|^2 - \frac{|\,(x - z|y)\,|^2}{\|y\|^2}.$$

It then follows that

$$(x - z|y) = 0.$$

This, of course, also holds for $y = 0$, so $x - z$ is orthogonal to Y, as promised. Note that we've already shown the existence of z if E is a Hilbert space and Y is closed, but we are not making that assumption here.

From the preceding results it follows that if F is a closed subspace of a Hilbert space E then every vector $x \in E$ can be written uniquely in the form

$$x = u + v$$

where $u \in F$ and v is orthogonal to F. Indeed, we take u to be the minimiser of $\|x - y\|$ over $y \in F$, which have already seen exists, and $v = x - u$, which we have already seen is orthogonal to F. In fact we only need to assume that F is a Hilbert space. E is allowed to be a non-complete inner product space, which can be useful when dealing with finite dimensional subspaces.

The set of vectors orthogonal to F is called the *orthogonal complement* of F and is denoted $F^\perp$. The decomposition above then shows that

$$E = F \oplus F^\perp.$$

Orthogonal complements are always closed. To see this, observe that if $y \in F$ then the set

$$\{x \in E\colon (x|y) = 0\}$$

is closed by the continuity of the inner product. Then

$$F^\perp = \cap_{y \in F} \{x \in E\colon (x|y) = 0\}$$

is the intersection of closed sets and hence also closed.

It is easy to see, from the fact that the orthogonality relation is symmetric, that

$$F^{\perp\perp} = F.$$

It follows that only closed spaces have orthogonal complements.

The *projection* from $P \in \mathcal{L}(E)$ defined by

$$Px = u$$

with x and u as above. Since $u \in F$ we have $Pu = u$ and hence $P^2 = P$. Since u and v are orthogonal we have

$$\|x\|^2 = \|u\|^2 + \|v\|^2$$

and hence

$$\|Px\|^2 = \|u\|^2 \leq \|x\|^2.$$

It follows that $\|P\| \leq 1$. On Assignment 4 we saw that $P^2 = P$ implies $\|P\| \geq 1$ unless $P = 0$ so

$$\|P\| = 1$$

unless $F = 0$.

Finally, if $x, y \in E$ then in addition to writing

$$x = u + v$$

where $u \in F$ and $v \in F^\perp$ we can write

$$y = w + z$$

where $w \in F$ and $z \in F^\perp$. We then have

$$\begin{aligned}(Px|y) &= (u|w + z) = (u|w) + (u|z) \\ &= (u|w) = (u + v|w) = (x|Py)\,.\end{aligned}$$

3.5 Orthogonal Sets

Suppose S is an orthonormal set for an inner product space E and $x \in E$ and $\delta > 0$. Let

$$S_\delta = \{u \in S\colon |\,(x|u)\,| \geq \delta\}$$

and suppose T is a finite subset of S_δ. By Bessel's inequality

$$\|x\|^2 \geq \sum_{u \in T} |\,(x|u)\,|^2 > n\delta^2$$

where n is the number of elements of T. So

$$n < \left(\frac{\|x\|}{\delta}\right)^2.$$

Since this holds for all finite subsets $T \subseteq S_\delta$ it follows that S_δ itself has at most that many elements. In particular, it is finite. If $(x|u) \neq 0$ then $u \subseteq S_{1/n}$ for some n, so it lies in

$$\cup_{n \in \mathbf{Z}^+} S_{1/n},$$

which is a countable union of finite sets and so is countable.

If $S = \{u_1, u_2, \ldots\}$ is a countable orthonormal set then

$$\sum_{j=1}^{\infty} \alpha_j u_j$$

is Cauchy if and only if

$$\sum_{j=1}^{\infty} |\alpha_j|^2$$

converges. This follows from one of the problems on Assignment 5. In a Hilbert space therefore

$$\sum_{j=1}^{\infty} \alpha_j u_j$$

converges if and only if

$$\sum_{j=1}^{\infty} |\alpha_j|^2$$

converges. It's a standard fact from Real Analysis that absolute convergence is unaffected by reordering the sequence, so the convergence of

$$\sum_{j=1}^{\infty} \alpha_j u_j$$

doesn't depend on the ordering in S. Combining this with the preceding paragraph, we can see how to give a meaning to

$$\sum_{u \in S} (x|u)\, u$$

for arbitrary, not necessarily countable, orthonormal sets S. We can eliminate all u for which $(x|u) = 0$, leaving countably many summands, which we may order however we choose. The sum will converge when E is a Hilbert space because the partial sums

$$\sum_{j=1}^{n} |\alpha_j|^2$$

are all bounded by $\|x\|$ and therefore form a bounded monotone sequence.

There's no guarantee that

$$\sum_{u \in S} (x|u)\, u,$$

if it converges, converges to x. The set S could, for example, be empty. In fact the sum converges to x if and only if

$$\|x\|^2 = \sum_{u \in S} |\, (x|u)\, |^2.$$

The remainder

$$x - \sum_{u \in S} (x|u)\, u$$

is orthogonal to all $u \in S$ and hence to the sum $\sum_{u \in S} (x|u)\, u$, from which we see that

$$\left\| x - \sum_{u \in S} (x|u)\, u \right\|^2 = \|x\|^2 - \left\| \sum_{u \in S} (x|u)\, u \right\|^2$$

and hence that

$$x = \sum_{u \in S} (x|u)\, u$$

if and only if

$$\|x\|^2 = \sum_{u \in S} |\, (x|u)\, |^2,$$

as claimed.

3.6 Orthonormal Bases

An *orthonormal basis* for an inner product space E is a set S such that for all $x \in E$

$$x = \sum_{u \in S} (x|u)\, u.$$

From the preceding section we see that this holds if and only if
$$\|x\|^2 = \sum_{u\in S} |\,(x|u)\,|^2.$$
An orthonormal subset of a Hilbert space is an orthonormal basis if and only if it is maximal. By Zorn's lemma any orthonormal subset is a subset of a maximal orthonormal set. In particular every Hilbert space has an orthonormal basis.

3.7 Continuous Linear Functionals

If E is an inner product space and $z \in E$ then
$$f(x) = (x|z)$$
defines a bounded linear transformation $f\colon E \to \mathbf{K}$. Linearity is clear from the properties of an inner product. Boundedness follows from
$$|f(x)| = |\,(x|z)\,| \le \|z\|\|x\|.$$
This shows that $\|z\|$ is a bound for f and hence that
$$\|f\| \le \|z\|.$$
To get the reverse inequality note that
$$|f(z)| = |\,(z|z)\,| = (z|z) = \|z\|^2.$$
We'll denote $\mathcal{L}(E, \mathbf{K})$ by E' and call it the *dual* of E. So we've associated to each $z \in E$ an $f \in E'$ with the same norm.

If E is a Hilbert space then all bounded linear functions from E to $\mathbf{K}$ are of this form. This is called the Riesz Representation Theorem[6] This is clear when $f = 0$, so assume $f \neq 0$ and let F be the null space of F, which must then be a proper subspace. It's closed because f was assumed to be continuous. We can therefore write
$$E = F \oplus F^{\perp}.$$
$F^{\perp} \neq \{0\}$ because $F \neq E$. Now f restricted to $F^{\perp}$ is an invertible linear transformation from $F^{\perp}$ to $\mathbf{K}$. Linearity follows from the linearity of f on E, the null space is zero by the definition of F and the range is all of $\mathbf{K}$ because otherwise it would have to be $\{0\}$, contradicting the assumption that $f \neq 0$. Let
$$g = (f|_{F^{\perp}})^{-1},$$
$$y = g(1)$$
and
$$z = \frac{1}{(y|y)} y.$$
Then if $x = u + v$ with $u \in F$ and $v \in F^{\perp}$,
$$v = g(f(v)) = g(f(v)1) = f(v)g(1) = f(v)y,$$
$$f(v) = f(x-u) = f(x) - f(u) = f(x)$$
and
$$(x|z) = (u + f(v)y|z) = (u|z) + f(v)\,(y|z) = f(x).$$

We can define a function h which takes $z \in E$ to the function $f \in E'$ given by $f(x) = (x|z)$. We've just shown that it's invertible if E is a Hilbert space. We've already seen that $\|h(z)\| = \|z\|$. It follows from the properties of inner products that
$$h(w+z) = h(w) + h(z)$$
and
$$h(\alpha z) = \overline{\alpha} h(z).$$
So h is linear if $\mathbf{K} = \mathbf{R}$, but is anti-linear if $\mathbf{K} = \mathbf{C}$.

3.8 Weak Convergence

Suppose that $(x_1, x_2, \ldots)$ is a sequence in a Hilbert space E and that $y \in E$. We say that the former converges weakly to the latter if, for all $z \in E$, we have
$$\lim_{n\to\infty} (x_n|z) = (y|z)\,.$$
Equivalently, we can say that it converges weakly if
$$f(x_n) \to f(x)$$
for all $f \in E'$, using the results of the previous section. We'll use the notation $x_n \rightharpoonup y$, which is common but not universal.

[6] Confusingly, several other results are also known by this name.

As an example of a sequence which converges weakly but not in the usual sense, consider the sequence $(e_1, e_2, \ldots) \in l^2$, where, as usual, e_j has its j'th entry equal to 1 and all others equal to 0. This converges weakly to 0, because if $z = (\zeta_1, \zeta_2, \ldots) \in l^2$ then

$$\lim_{n\to\infty} (e_n|z) = \lim_{n\to\infty} \overline{\zeta_n} = 0 = (0|z) .$$

If $x_n \to y$ then $x_n \rightharpoonup y$ and $\|x_n\| \to \|y\|$. The first of these statements follows from

$$|\,(x_n|z) - (y|z)\,| = |\,(x_n - y|z)\,| \le \|x_n - y\|\|z\|.$$

The second then follows from the fact that the norm is continuous. The converse is also true. If $x_n \rightharpoonup y$ and $\|x_n\| \to \|y\|$ then $x_n \to y$. To see this, note that

$$\|x_n - y\|^2 = 2\,\mathrm{Re}\,(x_n - y|y) + \left(\|x_n\|^2 - \|y\|^2\right) .$$

If $x_n \rightharpoonup y$ then $(x_n - y|y) \to 0$. If $\|x_n\| \to \|y\|$ then $\|x_n\|^2 - \|y\|^2 \to 0$. If both are true then $\|x_n - y\|^2 \to 0$ and hence $x_n \to y$.

Every bounded sequence has a weakly convergent subsequence. The proof of this is somewhat complicated. Suppose $(x_1, x_2, \ldots)$ is a bounded sequence, i.e. that

$$\|x_j\| \le \gamma$$

for all $j \in \mathbf{Z}^+$ and some $\gamma \in \mathbf{R}^+$. Use Gram-Schmidt, throwing away any x's which are linearly dependent on those earlier in the sequence to construct a countable orthonormal set $\{u_1, u_2, \ldots\}$ with the same span as $\{x_1, x_2, \ldots\}$. Let D be the set of vectors $w \in E$ for which $(w|u_j)$ is rational for each j and is zero for all but finitely many j. Then D is countable. Let $(y_1, y_2, \ldots)$ be a sequence which includes every element of D among its values. The sequence

$$((x_1|y_1)\,, \ldots (x_2|y_1)\,, \ldots)$$

is a bounded sequence in $\mathbf{K}$. By Bolzano-Weierstrass it has a convergent subsequence. In other words there is a subsequence of $(x_1, x_2, \ldots)$ whose inner products with y_1 converge. We can then look at the inner products of that subsequence with y_2 and select a subsequence whose inner products with y_2 converge. Continuing, we obtain a sequence of subsequences $x_{j,k}$ such that $(x_{1,1}, x_{1,2}, \ldots)$ is a subsequence of $(x_1, x_2, \ldots)$ and, for each $j \in \mathbf{Z}^+$, $(x_{j+1,1}, x_{j+1,2}, \ldots)$ is a subsequence of $x_{j,1}, x_{j,2}, \ldots)$ and, again for each $j \in \mathbf{Z}^+$,

$$\lim_{k\to\infty} (x_{j,k}|y_j) = \xi_j$$

exists. Set

$$z_n = x_{n,n}.$$

for each j the terms starting with the j'th one in the sequence $(u_1, u_2, \ldots)$ form a subsequence of $(z_{j,1}, z_{j,2}, \ldots)$ and so their inner products with y_j converge to the same limit. In other words

$$\lim_{n\to\infty} (z_n|y_j) = \xi_j$$

for each j. Let

$$f_n(x) = (x|z_n) .$$

Then

$$\|f_n\| = \|z_n\| \le \gamma,$$

since $(z_1, z_2, \ldots)$ is a subsequence of $(x_1, x_2, \ldots)$. If v belongs to the closure F of the span of $\{x_1, x_2, \ldots\}$, or equivalently of $\{u_1, u_2, \ldots\}$, then we can find a k with

$$\|v - y_k\| < \frac{\epsilon}{3\gamma}.$$

and hence

$$\|f_n(v) - f_n(y_k)\| < \frac{\epsilon}{\gamma}.$$

for all n. To find such a k we write

$$v = \sum_{j=1}^{\infty} (v|u_j)\, u_j$$

and choose an n such that

$$\|v - \sum_{j=1}^{n} (v|u_j)\, u_j\| < \frac{\epsilon}{6\gamma}$$

and then choose $\alpha_1, \ldots, \alpha_n \in \mathbf{Q}$ sufficiently close to $(v|u_1)\,, \ldots, (v|u_n)$ that

$$\|\sum_{j=1}^{n} (v|u_j)\, u_j - \sum_{j=1}^{n} \alpha_j u_j\| < \frac{\epsilon}{6\gamma}.$$

There is then, because of how the y's were defined, a k such that

$$y_k = \sum_{j=1}^{n} \alpha_j u_j.$$

For this k the sequence

$$(f_1(y_k), f_2(y_k), \ldots) = ((y_k|z_1), (y_k|z_2), \ldots)$$

is Cauchy, so there is an N such that if $m, n > N$ then

$$|f_m(y_k) - f_n(y_k)| < \frac{\epsilon}{3}.$$

Now

$$\begin{aligned} |f_m(v) - f_n(v)| \le {} & |f_m(v) - f_m(y_k)| \\ & + |f_m(y_k) - f_n(y_k)| \\ & + |f_n(v) - f_n(y_k)| < \epsilon \end{aligned}$$

so the sequence $(f_1(v), f_2(v), \ldots)$ has a limit, which we will call $f(v)$. It's easy to see that $f\colon F \to \mathbf{K}$ is linear. Also

$$f(v) = \lim_{n\to\infty} f_n(v)$$

implies

$$|f(v)| \le \limsup |f_n(v)| \le \gamma \|v\|$$

so $f \in F'$. We then extend f from F to

$$E = F \oplus F^{\perp}$$

by setting

$$f(w) = 0$$

for $w \in F^{\perp}$. By the Riesz Representation Theorem there is a $z \in E$ such that

$$f(x) = (x|z)$$

for all $x \in E$. Writing $x = v + w$ with $v \in F$ and $w \in F^{\perp}$ we see that

$$\begin{aligned} \lim_{n\to\infty} (x|z_n) &= \lim_{n\to\infty} (v|z_n) = \lim_{n\to\infty} f_n(v) \\ &= f(v) = f(x) = (x|z_n) \end{aligned}$$

and, since x was arbitrary,

$$z_n \rightharpoonup z.$$

4 Spectral Theory of Compact Symmetric Operators

4.1 Compact Operators

For simplicity we will refer to bounded linear transformations from here on as *operators*.[7] An operator A from a normed space E to a normed space F is called *compact* if the image in F of every bounded sequence in E has a convergent subsequence. If either E or F is finite-dimensional then this holds for all $A \in (E, F)$. For a less trivial example, set $E = F = C([a, b])$ and define A by

$$(Ax)(s) = \int_a^b k(s, t)x(t)\, dt$$

where k is a continuous $\mathcal{K}$-valued function on $[a, b] \times [a, b]$. The compactness of A then follows from the Arzelà-Ascoli Theorem.

It's straightforward to see that linear combinations of compact operators are compact. Also, the product of a compact operator with any operator, in either order, is compact. Suppose, for example, that $A\colon F \to G$ is a compact operator and $B\colon E \to F$ is an operator. If $(x_1, x_2, \ldots)$ is a bounded sequence in E then $(Bx_1, Bx_2, \ldots)$ is a bounded sequence in F and therefore $(ABx_1, ABx_2, \ldots)$, which is a sequence in G, has a convergent subsequence. If it's B which is compact then, if If $(x_1, x_2, \ldots)$ is a bounded sequence in E, the sequence $(Bx_1, Bx_2, \ldots)$ in F has a convergent subsequence. Because A is continuous the corresponding subsequence of $(ABx_1, ABx_2, \ldots)$ in G is also convergent.

If $\mathcal{A}_n \to B$ in $\mathcal{L}(E, F)$, F is a Banach space and A_n is compact for each n then B is compact. To see this, suppose that $(x_1, x_2, \ldots)$ is a bounded sequence in E, i.e. that there is a $\gamma \in \mathbf{R}^+$ such that

$$\|x_j\| \le \gamma$$

for all j. There's then a subsequence $(y_{1,1}, y_{1,2}, \ldots)$ of $(x_1, x_2, \ldots)$ such that $(A_1 y_{1,1}, A_1 y_{1,2}, \ldots)$ converges. This subsequence is still bounded, so we

[7] The term is often used for unbounded operators as well, but not for all unbounded operators. There is no commonly agreed definition of the term.

can extract a subsequence $(y_{2,1}, y_{2,2}, \ldots)$ such that $(A_2 y_{2,1}, A_2 y_{2,2}, \ldots)$ converges. Letting

$$z_n = y_{n,n}$$

we see that $(z_1, z_2, \ldots)$ is, for each j, equal from it's j'th term on to a subsequence of $(y_{j,1}, y_{j,2}, \ldots)$ and therefore that $(A_j z_1, A_j z_2, \ldots)$ is convergent. Then

$$Bz_j - Bz_k = A_n(z_j - z_k) + (A_n - B)z_k - (A_n - B)z_j.$$

If $\epsilon > 0$ then we choose n such that

$$\|A_n - B\| < \frac{\epsilon}{2\gamma + 1}$$

and hence

$$\|(A_n - B)z_j\| < \frac{\gamma\epsilon}{2\gamma + 1}$$

and

$$\|(A_n - B)z_k\| < \frac{\gamma\epsilon}{2\gamma + 1}$$

Because $(A_n z_1, A_n z_2, \ldots)$ is a Cauchy sequence there is an N such that if $j, k > N$ then

$$\|A_n(z_j - z_k)\| = \|A_n z_j - A_n z_k\| < \frac{\epsilon}{2\gamma + 1}.$$

Then, by the triangle inequality, we have

$$\|Bz_j - Bz_k\| < \epsilon.$$

So $(Bz_1, Bz_2, \ldots)$ is Cauchy and hence, by the completeness of F, is convergent. But $(z_1, z_2, \ldots)$ is a subsequence of $(x_1, x_2, \ldots)$, which was an arbitrary bounded sequence in E, so B is compact.

4.2 Symmetric Operators

We say that an operator $A \in \mathcal{L}(E)$, where E is an inner product space, is called *symmetric* if, for all $x, y \in E$,

$$(Ax|y) = (x|Ay)\,.$$

It follows immediately from the definition and the properties of inner products that sums of symmetric operators are symmetric, as are real multiples. If A and B are symmetric and $AB = BA$ then AB is symmetric because

$$(ABx|y) = (Bx|Ay) = (x|BAy) = (x|ABy)\,.$$

Without the assumption that A and B commute this need not be true.

We have already seen that orthogonal projections are symmetric. The operator $A \in \mathcal{L}(C([a, b]))$ defined by

$$(Ax)(s) = \int_a^b k(s, t)x(t)\,dt$$

is symmetric if

$$k(t, s) = \overline{k(s, t)}$$

for all $s, t \in [a, b]$, because

$$\begin{aligned}
(Ax|y) &= \int_a^b (Ax)(s)\overline{y(s)}\,ds \\
&= \int_a^b \left(\int_a^b k(s, t)x(t)\,dt\right)\overline{y(s)}\,ds \\
&= \int_a^b \int_a^b k(s, t)x(t)\overline{y(s)}\,dt\,ds \\
&= \int_a^b \int_a^b k(t, s)x(s)\overline{y(t)}\,ds\,dt \\
&= \int_a^b \int_a^b k(t, s)x(s)\overline{y(t)}\,dt\,ds
\end{aligned}$$

while

$$\begin{aligned}
(x|Ay) &= \int_a^b x(s)\overline{(Ay)(s)}\,ds \\
&= \int_a^b x(s)\overline{\int_a^b k(s, t)y(t)\,dt\,ds} \\
&= \int_a^b \int_a^b x(s)\overline{k(s, t)y(t)}\,dt\,ds \\
&= \int_a^b \int_a^b \overline{k(s, t)}x(s)\overline{y(t)}\,dt\,ds.
\end{aligned}$$

If $A \in \mathcal{L}(E)$ is symmetric and $x \in E$ then

$$\overline{(Ax|x)} = (x|Ax) = (Ax|x)\,,$$

so $(Ax|x)$ must be real. A symmetric operator is called *positive*[8] if

$$(Ax|x) \geq 0$$

[8]This terminology is standard, but is not ideal. Positive operators behave more like non-negative numbers than like positive numbers.

for all $x \in E$. Sums and non-negative real multiples of positive operators are positive.

It is easy to see that if A is positive then

$$q(x,y) = (Ax|y)$$

satisfies all the properties of an inner product except the last one. Since that property was not used in the proof of the Cauchy-Schwarz inequality it follows that

$$|q(x,y)|^2 \leq q(x,x)q(y,y),$$

i.e. that

$$|\,(Ax|y)\,|^2 \leq (Ax|x)\,(Ay|y)\,.$$

If A is a symmetric operator then

$$\|A\| = \sup_{\|x\|=1} |\,(Ax|x)\,| = \sup_{x\neq 0} \frac{|\,(Ax|x)\,|}{(x|x)}.$$

The second equation is easier. If $\|x\| = 1$ then $x \neq 0$ and

$$|\,(Ax|x)\,| = \frac{|\,(Ax|x)\,|}{(x|x)}$$

so, since the supremum over a larger set is at least as large as the supremum over a smaller set,

$$\sup_{\|x\|=1} |\,(Ax|x)\,| \leq \sup_{x\neq 0} \frac{|\,(Ax|x)\,|}{(x|x)}.$$

But if $x \neq 0$ and $y = x/\|x\|$ then $\|y\| = 1$ and

$$|\,(Ay|y)\,| = \frac{|\,(Ax|x)\,|}{(x|x)}.$$

so

$$\sup_{x\neq 0} \frac{|\,(Ax|y)\,|}{(x|x)} \leq \sup_{\|y\|=1} |\,(Ax|x)\,|.$$

Thus

$$\sup_{\|x\|=1} |\,(Ax|x)\,| = \sup_{x\neq 0} \frac{|\,(Ax|x)\,|}{(x|x)}.$$

From the Cauchy-Schwarz inequality we have

$$|\,(Ax|x)\,| \leq \|Ax\|\|x\| \leq \|A\|\|x\|^2 = \|A\|\,(x|x)$$

and hence

$$\frac{|\,(Ax|x)\,|}{(x|x)} \leq \|A\|$$

for $x \neq 0$ and, taking suprema,

$$\sup_{x\neq 0} \frac{|\,(Ax|x)\,|}{(x|x)} \leq \|A\|.$$

Let

$$w = \alpha x + \alpha^{-1} Ax, \quad z = \alpha x + \alpha^{-1} Ax.$$

where $\lambda \in \mathbf{R}^+$. Then

$$(Aw|w) - (Az|z) = 4\,(Ax|Ax) = 4\|Ax\|^2.$$

But

$$(Aw|w) \leq \sup_{x\neq 0} \frac{|\,(Ax|x)\,|}{(x|x)}\,(w|w)$$

and

$$(Az|z) \leq \sup_{x\neq 0} \frac{|\,(Ax|x)\,|}{(x|x)}\,(z|w)$$

so

$$4\|Ax\|^2 \leq \sup_{x\neq 0} \frac{|\,(Ax|x)\,|}{(x|x)} \left(\|w\|^2 + \|z\|^2\right).$$

By the Parallelogram identity

$$\|w\|^2 + \|z\|^2 = \leq 2\left(\alpha^2\|x\|^2 + \alpha^{-2}\|Ax\|^2\right)$$

We choose

$$\alpha = \sqrt{\frac{\|Ax\|}{\|x\|}}$$

so that the two terms in parentheses are both equal to $\|Ax\|\|x\|$ and

$$\|w\|^2 = \|z\|^2 = 4\|Ax\|\|x\|.$$

Combining everything,

$$4\|Ax\|^2 \leq 4 \sup_{x\neq 0} \frac{|\,(Ax|x)\,|}{(x|x)} \|Ax\|\|x\|$$

and hence

$$\|Ax\| \leq \sup_{x\neq 0} \frac{|\,(Ax|x)\,|}{(x|x)} \|x\|.$$

Thus the supremum on the right is a bound for A and hence

$$\|A\| \leq \sup_{x\neq 0} \frac{|\,(Ax|x)\,|}{(x|x)}.$$

We've already shown the reverse inequality, so

$$\|A\| = \sup_{x\neq 0} \frac{|\,(Ax|x)\,|}{(x|x)}.$$

4.3 Eigenspace Decomposition

We say that λ is an *eigenvalue* and v an *eigenvector* of an operator A if

$$Av = \lambda v$$

with $v \neq 0$. Suppose now that A is symmetric. If v is an eigenvector with eigenvalue λ and w is an eigenvector with eigenvalue $\mu 0$ then

$$\begin{aligned} \lambda\,(v|w) &= (\lambda v|w) = (Av|w) = (v|Aw) \\ &= (v|\mu w) = \overline{\mu}\,(v|w)\,, \end{aligned}$$

so either

$$\lambda = \overline{\mu}$$

or

$$(v|w) = 0.$$

In the case $\lambda = \mu$, $v = w$ we can't have $(v|w) = 0$ so $\lambda = \overline{\mu}$. In other words, the eigenvalues of a symmetric operator are all real. If $\lambda \neq \mu$ then, since both are real, $\lambda \neq \overline{\mu}$ and hence $(v|w) = 0$. In other words, eigenvectors corresponding to distinct eigenvectors are orthogonal. If A is positive then

$$\lambda\,(v|v) = (\lambda v|v) = (Av|v) \geq 0$$

so

$$\lambda \geq 0.$$

In other words the eigenvalues of a positive operator are non-negative.

Assume for the moment that A is a non-zero compact symmetric operator. We saw at the end of the last section that

$$\|A\| = \sup_{\|x\|=1} |\,(Ax|x)\,|$$

so there is a sequence $(x_1, x_2, \ldots)$ with

$$\|x_n\| = 1$$

for each n and

$$\lim_{n\to\infty} |\,(Ax_n|x_n)\,| = \|A\| \neq 0.$$

It follows that, except possibly for finitely many n, the quantity $(Ax_n|x_n)$ is non-zero. It's always real, so there are infinitely many n for which it's positive or infinitely many for which it's negative. There could be both, of course, but in any case we have a subsequence $(y_1, y_2, \ldots)$ of $(x_1, x_2, \ldots)$ for which

$$\lim_{n\to\infty} (Ay_n|y_n) = \lambda.$$

and

$$|\lambda| = \|A\|.$$

Taking limits in the inequalities

$$0 \leq \|Ay_n - \lambda y_n\|^2 = \|Ay_n\|^2 - 2\lambda\,(Ay_n|y_n) + \lambda^2$$

we see that

$$\lim_{n\to\infty} (Ay_n - \lambda y_n) = 0.$$

The sequence $(y_1, y_2, \ldots)$ is bounded, because $\|y_n\| = 1$ for all n, so by the compactness of A we can find a subsequence $(z_1, z_2, \ldots)$ such that $(Az_1, Az_2, \ldots)$ is convergent:

$$\lim_{n\to\infty} z_n = u.$$

By the continuity of the norm we have

$$\|u\| = 1.$$

Also

$$\lim_{n\to\infty} (Az_n - \lambda z_n) = 0,$$

so

$$Au - \lambda u = 0.$$

So u is an eigenvector with eigenvalue λ.

The preceding construction can be repeated. If A_1 is a compact symmetric operator on an infinite dimensional Hilbert space E_1 then set

$$E_1 = E, \quad A_1 = A.$$

After finding an eigenvector u_1 of norm 1 with eigenvalue λ_1,

$$\|u_1\| = 1, \quad Au_1 = \lambda_2 u_2,$$

we can form the subspace

$$E_2 = \{x \in E_1 \colon (x|u_1) = 0\}$$

and the operator A_2 on E_2 obtained by restricting A_1 to E_2. This makes sense, since if $x \in E_2$ then

$$(A_1 x | u_1) = (x | A u_1) = (x | \lambda_1 u_1) = \overline{\lambda}\,(x | u_1) = 0,$$

so $A_1 x \in E_2$. A_2 is also compact and symmetric. Indeed if $(x_1, x_2, \ldots)$ is a bounded sequence in E_2 then it's a bounded sequence in E_1 and hence has a subsequence $(y_1, y_2, \ldots)$ such that $(A_1 y_1, A_1 y_2, \ldots)$ is convergent in E_1. But each $y_n \in E_2$, A_2 is the restriction to E_2 of A_1, A_2 multiplied by anything in E_2 is, as we've seen, in E_2 and E_2 is closed, so $(A_2 y_1, A_2 y_2, \ldots)$ is convergent in E_2. That establishes compactness. Symmetry is easier. If $x, y \in E_2$ then

$$(A_2 x | y) = (A_1 x | y) = (x | A_1 y) = (x | A_2 y)\,.$$

We can therefore find a u_2 in E_2, of norm 1, which is an eigenvector of A_2 with eigenvalue λ_2. In other words,

$$\|u_2\| = 1, \quad A_2 u_2 = \lambda_2 u_2.$$

Since A_2 is the restriction of $A = A_1$ we could equally well write

$$\|u_2\| = 1, \quad A u_2 = \lambda_2 u_2.$$

Also since $u_2 \in E_2$ we have

$$(u_1 | u_2) = 0.$$

We can then define

$$E_3 = \{x \in E_2 \colon (x | u_2) = 0\}$$

and let A_3 be the restriction to E_3 of A_2 or, equivalently, of A. We then find u_3 and λ_3 such that

$$\|u_3\| = 1, \quad A u_3 = \lambda_3 u_3,$$

$$(u_1 | u_3) = 0, \quad (u_2 | u_3) = 0.$$

The process described above will fail if at some stage $A_n = 0$, which may or may not happen. If this happens then we construct u_n, u_{n+1}, …by taking an infinite orthonormal set in E_n. Since $A_n = 0$ the elements of this set are all eigenvectors of eigenvalue 0. So even if this happens we still obtain an orthonormal set $\{u_1, u_2, \ldots\}$ of eigenvectors and a corresponding set of eigenvalues $\{\lambda_1, \lambda_2, \ldots\}$.

$$A u_j = \lambda_j, \quad (u_j | u_k) = \begin{cases} 1 & \text{if } j = k, \\ 0 & \text{if } j \neq k. \end{cases}$$

Let $\epsilon > 0$. If there are infinitely many j with $|\lambda_j| \geq \epsilon$ then, selecting only the corresponding u_j, we can form a bounded sequence $(v_1, v_2, \ldots)$. We then have

$$\|A v_j - A v_k\| \geq \sqrt{2}\epsilon$$

for all j, k, so the sequence $(A v_1, A v_2, \ldots)$ is not Cauchy and hence not convergent. This contradicts the compactness of A, so the assumption that there are infinitely many such j. Thus there are only finitely many j with $|\lambda_j| \geq \epsilon$.

$$N = \max\{j \in \mathbf{Z}^+ \colon |\lambda_j| \geq \epsilon\}.$$

Then

$$j > N$$

implies

$$|\lambda_j - 0| < \epsilon.$$

There is such an N for each $\epsilon > 0$. In other words,

$$\lim_{n \to \infty} \lambda_n = 0.$$

The orthonormal set $\{u_1, u_2, \ldots\}$ could still fail to be an orthonormal basis. Indeed it must fail to be an orthonormal basis if our Hilbert space is so large that it doesn't have a countable orthonormal basis. We can show though that if

$$(x | u_j) = 0$$

for all j then

$$A x = 0.$$

Indeed $(x | u_j) = 0$ for all j implies

$$x \in E_j$$

for all j and hence

$$\begin{aligned} (Ax | x) = (A_j x | x) &\leq \|A_j x\| \|x\| \\ &= \|A_j\| \|x\|^2 = |\lambda_j| \|x\|^2. \end{aligned}$$

Taking limits as $j \to \infty$ we find that

$$(Ax|x) = 0.$$

We apply the analogue

$$(Ax|y) \leq \sqrt{(Ax|x)}\sqrt{(Ay|y)}$$

of the Cauchy-Schwarz Inequality derived earlier with $y = Ax$ to get

$$\|Ax\|^2 \leq (Ax|y) \leq \sqrt{(Ax|x)}\sqrt{(Ay|y)} = 0$$

and hence $\|Ax\|^2 = 0$ and

$$Ax = 0.$$

So even if $\{u_1, u_2, \ldots\}$ isn't maximal we can extend it to an orthonormal set consisting entirely of eigenvectors of A, at most countably many of which have non-zero eigenvalues. We thus always obtain an orthonormal set S and a function $\lambda\colon S \to \mathbf{R}$ such that for all $x \in E$ we have

$$x = \sum_{u \in S} (x|u)\, u, \quad Ax = \sum_{u \in S} (x|u)\, \lambda(u) u.$$

Another way to express the decomposition above is to let Λ be the set of all $\mu \in \mathbf{R}$ such that $\mu = \lambda(u)$ for some $u \in S$ and

$$S_\mu = \{u \in S\colon \lambda(u) = \mu\}, \quad P_\mu x = \sum_{u \in S_\mu} (x|u)\, u.$$

where $\mu \in \Lambda$. The P's are called spectral projections. We then have the following relations:

$$P_\mu P_\nu = P_\nu P_\mu = \begin{cases} P_\mu & \text{if } \mu = \nu, \\ 0 & \text{if } \mu \neq \nu, \end{cases}$$

$$\sum_{\mu \in \Lambda} P_\mu x = x$$

and

$$AP_\mu = P_\mu A = \mu P_\mu.$$

for all $\mu, \nu \in \mathbf{\Lambda}$ and $x \in E$. It's tempting to say that since $\sum_{\mu \in \mathbf{R}} P_\mu x = x$ for all x we must have

$$\sum_{\mu \in \mathbf{R}} P_\mu = I,$$

but this is not true, at least with the meaning to which we've given infinite sums up until now.

4.4 The Equation $(\mu I - A)x = y$

Suppose A is a compact symmetric operator on a Hilbert space E and let Λ be it's set of eigenvalues. so that there is an orthonormal set S and a function $\lambda\colon S \to \Lambda$ such that

$$Au = \lambda(u)u$$

for all $u \in S$. If $x \in S$ then

$$x = \sum_{u \in S} (x|u)\, u$$

and

$$Ax = \sum_{u \in S} (x|u)\, \lambda(u) u.$$

If

$$y = (\mu I - A)x = \mu x - Ax$$

then, taking the inner product with $v \in S$,

$$(y|v) = (\mu - \lambda(v))\, (x|v)\, .$$

If $\mu \notin \Lambda$ then we can solve this to obtain

$$(x|v) = \frac{(y|v)}{\mu - \lambda(v)}$$

and hence

$$x = \sum_{v \in S} \frac{(y|v)}{\mu - \lambda(v)} v.$$

So far we have assumed the existence of x. But if $\mu \notin \Lambda$ as we've already assumed, and also $\mu \neq 0$ then we can show that

$$\sum_{v \in S} \frac{(y|v)}{\mu - \lambda(v)} v.$$

converges and that its limit solves

$$(\mu I - A)x = y.$$

The convergence follows from

$$\sum_{v \in S} \left| \frac{(y|v)}{\mu - \lambda(v)} \right|^2 \leq \gamma \sum_{v \in S} |\, (y|v)\, |^2 = \gamma \|y\|^2$$

where

$$\gamma = \max_{v\in S} \frac{1}{|\mu-\lambda(v)|^2}. = \max \lambda \in \Lambda \frac{1}{|\mu-\lambda|^2}.$$

The maximum is finite because

$$\overline{\Lambda} \subseteq \Lambda \cup \{0\}$$

is compact. The bar here denotes closure rather than complex conjugation. If we call the limit x then

$$x = \sum_{v\in S} \frac{1}{\mu-\lambda(v)} (y|v)\, v,$$

$$\mu x = \sum_{v\in S} \frac{\mu}{\mu-\lambda v} (y|v)\, v,$$

and

$$Ax = \sum_{v\in S} \frac{\lambda(v)}{\mu-\lambda v} (y|v)\, v,$$

so

$$\mu x - Ax \sum_{v\in S} (y|v)\, v = y.$$

The equation

$$(\mu I - A)x = y$$

therefore has a unique solution x for all $y \in E$ provided that $\mu \notin \Lambda$ and $\mu \neq 0$.

If $\mu \in \Lambda$ then the situation is more complicated. It follows from

$$(y|v) = (\mu - \lambda(v))\,(x|v)$$

that $(y|v) = 0$ for all v such that $\lambda(v) = \mu$. In terms of the spectral projection

$$P_\mu y = 0.$$

This is therefore a necessary condition for the equation

$$(\mu I - A)x = y$$

to have a solution. If $\mu \neq 0$ then it is also a sufficient condition, since

$$x = \sum_{v\in S-S_\mu} \frac{1}{\mu-\lambda(v)} (y|v)\, v$$

is a a solution. There is no uniqueness however, since

$$w + \sum_{v\in S-S_\mu} \frac{1}{\mu-\lambda(v)} (y|v)\, v$$

is also a solution for any eigenvector w of A with eigenvalue μ.

4.5 Sturm-Liouville Problems

Many problems in mathematical physics can be reduced to solving ordinary differential equations of the form

$$p(t)x''(t) + p'(t)x'(t) + q(t)x(t) + \mu r(t)x(t) = y(t)$$

subject to certain boundary conditions. In this section all constants and functions will be assumed to be real. For simplicity we will restrict our attention here to problems on a compact interval $[a,b]$ where p, p', q, r and y are all continuous and p and r are positive. Also all constants and functions in this section will be real. The boundary conditions we'll consider are of the form

$$\alpha_0 x(a) + \alpha_1 x'(a) = 0, \quad \beta_0 x(b) + \beta_1 x'(b) = 0$$

with (α_0, α_1) and (β_0, β_1) non-zero. Boundary value problems of this type are called Sturm-Liouville problems.

Integration by parts shows that if x, z are twice differentiable and

$$p(t)x''(t) + p'(t)x'(t) + q(t)x(t) + \mu r(t)x(t) = 0$$

then

$$\begin{aligned}\int_a^b &[p(t)x'(t)z'(t) - q(t)x(t)z(t)]\; dt \\ &= p(b)x'(b)z(b) - p(a)x'(a)z(a) \\ &\quad + \mu \int_a^b r(t)x(t)z(t)\, dt.\end{aligned}$$

If also

$$p(t)z''(t) + p'(t)z'(t) + q(t)z(t) + \nu r(t)z(t) = 0$$

then, switching the roles of x and z and of μ and ν,

$$\int_a^b [p(t)x'(t)z'(t) - q(t)x(t)z(t)]\ dt$$
$$= p(b)x(b)z'(b) - p(a)x(a)z'(a)$$
$$+ \nu \int_a^b r(t)x(t)z(t)\, dt.$$

Combining these,

$$(\mu - \nu) \int_a^b r(t)x(t)z(t)\, dt = w(b) - w(a)$$

where

$$w(t) = p(t)\left[x'(t)z(t) - x(t)z'(t)\right]$$

If

$$\alpha_0 x(a) + \alpha_1 x'(a) = 0, \quad \alpha_0 z(a) + \alpha_1 z'(a) = 0$$

then $w(a) = 0$ and if

$$\beta_0 x(b) + \beta_1 x'(b) = 0, \quad \beta_0 z(b) + \beta_1 z'(b) = 0$$

then $w(b) = 0$. If both are true then

$$(\mu - \nu) \int_a^b r(t)x(t)z(t)\, dt = 0.$$

So if $\mu \neq \nu$ then

$$\int_a^b r(t)x(t)z(t)\, dt = 0.$$

In other words x and z are orthogonal elements of the Hilbert space E of measurable functions x for which

$$\int_a^b r(t)x(t)^2\, dt < \infty.$$

If $r = 1$ then $E = L^2([a, b])$. As in that special case we need to take the quotient by the space of functions which vanish almost everywhere in order to get norm rather than a semi-norm. The corresponding inner product is

$$(x|z) = \int_a^b r(t)x(t)z(t)\, dt.$$

We can try to view the equation

$$p(t)x''(t) + p'(t)x'(t) + q(t)x(t) + \mu r(t)x(t) = 0$$

as an eigenvector-eigenvalue equation

$$Lx = \mu x$$

with

$$L = -\frac{1}{r(t)} \left(\frac{d}{dt} p(t) \frac{d}{dt} + q(t) \right).$$

This would explain why eigenvectors corresponding to distinct eigenvalues are orthogonal. But we quickly run into problems. This differential operator is not well defined on all of E. Worse still, even on the subset where it is well defined it is not bounded, and hence has no chance of being compact. It is possible, and useful, to develop a version of spectral theory for densely defined, unbounded operators, but this is difficult. Fortunately there is a way around this for Sturm-Liouville problems.

Choose φ_0 and φ_1 such that

$$\alpha_0 \varphi_0 + \alpha_1 \varphi_1 = 0,$$

but $(\varphi_0, \varphi_1) \neq (0, 0)$. By the existence and uniqueness theorem for ordinary differential equations there is a unique twice continuously differentiable function $u\colon \mathbf{R} \times [a, b] \to \mathbf{R}$ such that

$$p(t)u''(\mu, t) + p'(t)u'(\mu, t) + q(t)u(\mu, t) + \mu r(t)u(\mu, t) = 0,$$

$$u(\mu, a) = \varphi_0, \quad u'(\mu, a) = \varphi_1.$$

Primes here denote t derivatives. Similarly, if

$$\beta_0 \psi_0 + \beta_1 \psi_1 = 0,$$

but $(\psi_0, \psi_1) \neq (0, 0)$, then there's a unique solution v to

$$p(t)v''(\mu, t) + p'(t)v'(\mu, t) + q(t)v(\mu, t) + \mu r(t)v(\mu, t) = 0,$$

$$v(\mu, b) = \psi_0, \quad v'(\mu, b) = \psi_1.$$

We then define

$$w(\mu, t) = p(t)[u(t)v'(t) - u'(t)v(t)]$$

and note that

$$w'(\mu, t) = 0.$$

Abusing notation slightly, we write

$$w(\mu, t) = w(\mu).$$

Let

$$k(\mu, s, t) = \begin{cases} u(\mu, s)v(\mu, t) & \text{if } a \le s \le t \le b, \\ v(\mu, s)u(\mu, t) & \text{if } a \le t \le s \le b. \end{cases}$$

If

$$z(\mu, s) = \int_a^b k(\mu, s, t)y(t)\, dt$$

then direct calculation shows that

$$\begin{aligned} z(\mu, a) &= \int_a^b u(\mu, a)v(\mu, t)y(t)\, dt \\ &= \alpha_0 \int_a^b v(\mu, t)y(t)\, dt, \end{aligned}$$

$$\begin{aligned} z(\mu, b) &= \int_a^b v(\mu, b)u(\mu, t)y(t)\, dt \\ &= \beta_0 \int_a^b u(\mu, t)y(t)\, dt, \end{aligned}$$

$$\begin{aligned} z'(\mu, s) = &\int_a^s v'(\mu, s)u(\mu, t)y(t)\, dt \\ &+ \int_s^b u'(\mu, s)v(\mu, t)y(t)\, dt, \end{aligned}$$

$$\begin{aligned} z'(\mu, a) &= \int_a^b u'(\mu, a)v(\mu, t)y(t)\, dt \\ &= \alpha_1 \int_a^b v(\mu, t)y(t)\, dt, \end{aligned}$$

$$\begin{aligned} z'(\mu, b) &= \int_a^b v'(\mu, b)u(\mu, t)y(t)\, dt \\ &= \beta_1 \int_a^b u(\mu, t)y(t)\, dt, \end{aligned}$$

$$\alpha_0 z(\mu, a) + \alpha_1 z'(\mu, a) = 0,$$

$$\beta_0 z(\mu, b) + \beta_1 z'(\mu, b) = 0,$$

$$\begin{aligned} z''(\mu, s) = &\int_a^s v''(\mu, s)u(\mu, t)y(t)\, dt \\ &+ \int_s^b u''(\mu, s)v(\mu, t)y(t)\, dt, \\ &+ [u(\mu, s)v'(\mu, s) - u'(\mu, s)v(\mu, s)]y(s) \\ = &\int_a^s v''(\mu, s)u(\mu, t)y(t)\, dt \\ &+ \int_s^b u''(\mu, s)v(\mu, t)y(t)\, dt + \frac{w(\mu)}{p(s)} y(s), \end{aligned}$$

$$\begin{aligned} &p(s)z''(\mu, s) + p'(s)z'(\mu, s) + q(s)z(\mu, s) + \mu r(s)z(\mu, s) \\ &\quad = w(\mu)y(s). \end{aligned}$$

For each $\mu \in \mathbf{R}$ either $w(\mu) = 0$ or $w(\mu) \neq 0$. If $w(\mu) \neq 0$ then preceding calculation shows us that

$$x(\mu, s) = \frac{z(\mu, s)}{w(\mu)}$$

is a solution to

$$px'' + p'x' + qx + \mu rx = y,$$

$$\alpha_0 z(\mu, a) + \alpha_1 z'(\mu, a) = 0,$$

$$\beta_0 z(\mu, b) + \beta_1 z'(\mu, b) = 0.$$

If $w(\mu) = 0$ then u and v are linearly dependent, since

$$u(a)v'(a) - u'(a)v(a) = \frac{w(\mu, a)}{p(a)} = 0$$

which means that $(v(a), v'(a))$ is a multiple of $(u(a), u'(a)) = (\alpha_0, \alpha_1)$ and, by the uniqueness part of the existence and uniqueness theorem for the initial value problem, v is the same multiple of u. So there is then a non-zero solution to

$$px'' + p'x' + qx + \mu rx = 0,$$

$$\alpha_0 z(\mu, a) + \alpha_1 z'(\mu, a) = 0,$$

$$\beta_0 z(\mu, b) + \beta_1 z'(\mu, b) = 0.$$

Define $A_\mu \in \mathcal{L}(E)$ by

$$(A_\mu x)(s) = \int_a^b k(\mu, s, t)r(t)x(t)\, dt.$$

A_μ is symmetric because

$$\begin{aligned}((A_\mu)x|y) &= \int_a^b r(s)(A_\mu x)(s)y(s)\,ds \\ &= \int_a^b \int_a^b r(s)k(\mu,s,t)r(t)x(t)y(s)\,dt\,ds \\ &= \int_a^b \int_a^b r(s)k(\mu,s,t)r(t)x(t)y(s)\,ds\,dt,\end{aligned}$$

$$\begin{aligned}(x|(A_\mu)y) &= \int_a^b r(s)x(s)(A_\mu y)(s)\,ds \\ &= \int_a^b \int_a^b r(s)k(\mu,s,t)r(t)y(t)x(s)\,dt\,ds \\ &= \int_a^b \int_a^b r(s)k(\mu,t,s)r(t)y(s)x(t)\,dt\,ds\end{aligned}$$

and $k(\mu,s,t) = k(\mu,t,s)$. Also A_μ is compact because of the Arzelà-Ascoli Theorem.

It can be shown that w is not identically zero. We can therefore choose a μ for which $w(\mu) \neq 0$. Let x be an eigenvector of A_μ with eigenvalue λ,

$$y(t) = r(t)x(t)$$

$$\begin{aligned}z(s) &= \int_a^b k(\mu,s,t)y(t)\,dt \\ &= \int_a^b k(\mu,s,t)r(t)x(t)\,dt. \\ &= (A_\mu x)(s) = \lambda x(s).\end{aligned}$$

By the calculation we did earlier,

$$\begin{aligned}&p(s)z''(s) + p'(s)z'(s) + q(s)z(s) + \mu r(s)z(s) \\ &\quad = w(\mu)y(s),\end{aligned}$$

or

$$p(s)x''(s) + p'(s)x'(s) + q(s)x(s) + \nu r(s)x(s) = 0$$

where

$$\nu = \mu - w(\mu)/\lambda.$$

Constructing an orthonormal set $\{x_1, x_2, \ldots\}$ of eigenvectors of A_μ with eigenvalues $\lambda_1, \lambda_2, \ldots$ with

$$\lim_{n\to\infty} \lambda_n = 0,$$

as in Section 4.3, we have

$$p(s)x_n''(s) + p'(s)x_n'(s) + q(s)x_n(s) + \nu_n r(s)x_n(s) = 0$$

where

$$\lim_{n\to\infty} |\nu_n| = \infty.$$

In fact it can be shown that

$$\lim_{n\to\infty} \nu_n = +\infty.$$

As in Section 4.3 we have to worry about the possibility that our orthonormal set is not complete. We saw there though that this can only happen if the null space of A_μ is non-trivial. We've seen that

$$\begin{aligned}&p(s)z''(s) + p'(s)z'(s) + q(s)z(s) + \mu r(s)z(s) \\ &\quad = w(\mu)x(s)\end{aligned}$$

where $z = A_\mu x$. If x belongs to the null space of A_μ then the left hand side is zero, so the right hand side is as well. But μ was chosen such that $w(\mu) \neq 0$, so $x = 0$. The null space of A_μ is therefore trivial.

5 The Main Theorems of Functional Analysis

5.1 Hahn-Banach

Suppose E is a vector space over $\mathbf{K}$, p is a semi-norm on E, F is a subspace of E, f is a linear transformation from F to $\mathbf{K}$ and, for all $x \in F$,

$$|f(x)| \leq p(x).$$

In that case there is linear transformation g from E to $\mathbf{K}$ such that

$$g(x) = f(x)$$

for all $x \in F$ and

$$|g(x)| \leq p(x)$$

for all $x \in E$. This result is known as the Hahn-Banach Theorem, or rather as the semi-norm version of the Hahn-Banach Theorem. There is also a norm version, which we will derive from the semi-norm version once we have proved it. The Hahn-Banach Theorem should be thought of as an extension theorem because the equation $g(x) = f(x)$ says that g is an extension of f from F to E.

Consider the set P of pairs (H, h) where H is a subspace of E, F is a subspace of H, h is linear on H

$$h(x) = f(x)$$

for all $x \in F$ and

$$|h(x)| \leq p(x)$$

for all $x \in H$. On P we impose a partial order by saying

$$(H_1, h_1) \leq (H_2, h_2)$$

if $H_{\subseteq} H_2$ and

$$h_1(x) = h_2(x)$$

for all $x \in H_1$. Any totally ordered subset T of P has a bound, namely (G, g) where

$$G = \cup_{(H,h)\in T} H$$

and

$$g(x) = h(x)$$

if $x \in H$, $(H, h) \in T$. This last equation defines $g(x)$ unambiguously because if $x \in H_1$ and $x \in H_2$ then either $H_1 \subseteq H_2$ or $H_2 \subseteq H_1$ and then $h_1(x) = h_2(x)$. The hypotheses of Zorn's lemma are therefore satisfied, so there is a maximal $(H, h) \in P$.

For the moment, assume $\mathbf{K} = \mathbf{R}$. If $u, v \in H$ and $z \in E$ then

$$\begin{aligned} h(u) - h(v) &= h(u - v) \leq p(u - v) \\ &= p(u + z - v - z) \\ &\leq p(u + z) + p(v + z). \end{aligned}$$

Equivalently,

$$-p(v + z) - h(v) \leq p(u + z) - h(u).$$

Taking the supremum over $v \in H$ and the infimum over $u \in H$ gives the inequality

$$\sup_{y\in H} \left[-p(y + z) - h(y)\right] \leq \inf_{y\in H} \left[p(y + z) - h(y)\right].$$

Let ξ be the average of the left and right hand sides of this inequality. Then

$$-p(y + z) - h(y) \leq \xi \leq p(y + z) - h(y)$$

for all $y \in H$. If $\alpha > 0$ then $y/\alpha \in H$ and, replacing y by y/α in

$$\xi \leq p(y + z) - h(y)$$

gives

$$\xi \leq p(y/\alpha + z) - h(y/\alpha)$$

and hence

$$h(y) + \alpha\xi \leq p(y + \alpha z).$$

Similarly, if $\alpha < 0$ then, replacing y by y/α in

$$-p(y + z) - h(y) \leq \xi$$

gives

$$-p(y/\alpha + z) - h(y/\alpha) \leq \xi,$$

and

$$-p(y + \alpha z) \leq h(y) + \alpha\xi$$

For any $\alpha \neq 0$ and $y \in H$ then

$$|h(y) + \alpha\xi| \leq p(y + \alpha z)$$

This also holds if $\alpha = 0$ though, since $h(y) \leq p(y)$ for all $y \in H$.

Let K be the space spanned by H and z. In other words, $x \in G$ if and only if

$$x = y + \alpha z$$

for some $y \in H$ and $\alpha \in \mathbf{R}$. This y and α are uniquely determined if $z \notin K$, so it makes sense to define

$$k(x) = h(y) + \alpha\xi.$$

But then

$$k(x) = h(y)$$

and

$$|k(x)| \leq p(x)$$

for all $x \in K$. Also $H \subseteq K$. In other words,

$$(H, h) \leq (K, k).$$

This contradicts the maximality of (H, h), so there can be no $z \in E - H$. In other words, $H = E$. This proves the theorem in the case $\mathbf{K} = \mathbf{R}$.

Suppose now that $\mathbf{K} = \mathbf{C}$. Since $\mathbf{R} \subseteq \mathbf{C}$ we can consider F as a vector space over $\mathbf{R}$. We write

$$f_1(x) = \operatorname{Re} f(x)$$

and

$$f_2(x) = \operatorname{Im} f(x)$$

so that

$$f(x) = f_1(x) + i f_2(x).$$

Then f_1 and f_2 are linear transformations from E, considered as a real vector space, to $\mathbf{R}$. We have

$$|f_1(x)| \le |f(x)| \le p(x)$$

for all $x \in F$ and so, by the real version of the theorem, which we have just proved, there is a linear $g_1 \colon E \to \mathbf{R}$ such that

$$g_1(x) = f_1(x)$$

for all $x \in F$ and

$$g_1(x) \le p(x)$$

for all $x \in E$. Define $g \colon E \to \mathbf{C}$ by

$$g(x) = g_1(x) - i g_1(ix).$$

Then

$$g(x) = f(x)$$

for all $x \in F$. Also, using the polar decomposition

$$g(x) = \rho e^{i\theta},$$

we have that $e^{-i\theta} g(x)$ is real and hence

$$\begin{aligned} |g(x)| &= |e^{-i\theta} g(x)| \\ &= |g(e^{-i\theta} x)| \\ &= |g_1(e^{-i\theta} x)| \\ &\le |p(e^{-i\theta} x)| \\ &\le |p(x)|. \end{aligned}$$

This proves the complex version of the theorem.

The main point of the semi-norm version of the Hahn-Banach Theorem is to prove the norm version, which says that if F is a subspace of a normed space E and $f \in F' = \mathcal{L}(F, \mathbf{K})$ then there is a $g \in E' = \mathcal{L}(E, \mathbf{K})$ with

$$g(x) = f(x)$$

for all $x \in F$ and

$$\|g\| = \|f\|.$$

To prove this we apply the semi-norm version of the theorem with

$$p(x) = \|f\| \|x\|.$$

This gives an extension g of f from F to E satisfying

$$\|g(x)\| \le \|f\| \|x\|$$

from which is follows that $g \in E'$ and

$$\|g\| \le \|f\|.$$

But if $x \in F$ then

$$|f(x)| = |g(x)| \le \|g\| \|x\|$$

so

$$\|f\| \le \|g\|$$

and hence

$$\|g\| = \|f\|.$$

As a special case, note that if F is finite dimensional then any linear transformation $f \colon F \to \mathbf{K}$ is bounded and so can be extended to a bounded transformation on E. In particular this shows that E' is non-trivial, i.e. not $\{0\}$, if E is non-trivial.

5.2 The Baire Category Theorem

The diameter of a metric space X with metric d is defined to be

$$\delta(M) = \sup_{x,y \in X} d(x, y).$$

This applies to any subset of a metric space, since subsets of metric spaces are themselves metric spaces. Suppose that X is a complete metric space and $S_1 \supseteq S_2 \supseteq \cdots$ are non-empty closed subsets of X with $\delta(S_n) \to 0$. Then

$$\cap_{n \in \mathbf{Z}^+} S_n$$

contains exactly one point. To see this, we choose $x_j \in S_j$ for each j. If $\epsilon > 0$ then there is an N such that if $j > N$ then $\delta(S_j) < \epsilon$. If $j, k > N$ then

$$x_j, x_k \in S_{\max(j,k)}$$

and hence

$$d(x_j, x_k) \leq \delta(S_{\max(j,k)}) < \epsilon.$$

So $(x_1, x_2, \ldots)$ is a Cauchy sequence and hence

$$y = \lim_{n\to\infty} x_n$$

exists. For any m we then have

$$\lim_{n\to\infty} x_{m+n} = y.$$

But $x_{m+n} \in S_{m+n} \subseteq S_m$ and S_m is closed, so

$$y \in S_m.$$

Since this is true for all m we have

$$y \in \cap_{m\in\mathbf{Z}^+} S_m$$

If $z \in \cap_{m\in\mathbf{Z}^+} S_m$ as well then

$$d(y, z) \leq \delta(S_m)$$

for all m and hence $d(y, z) = 0$ and therefore

$$y = z.$$

The Baire Category Theorem says that if a complete metric space X is a countable union

$$X = \cup_{n\in\mathbf{Z}^+} S_n$$

of closed sets S_n then at least one of the S_n contains an open ball. The proof is by contradiction.

We begin by establishing a lemma that if $S \subset X$ is closed, $x_1 \in S$ and $r_1 > 0$ then either the open ball B_1 of radius r_1 about x_1 is contained in S or there's an $x_2 \in X - S$ and an $r_2 > 0$ such that the closed ball B_2 of radius r_2 about x_2 is contained in $X - S$ and $B_2 \subseteq B_1$. The lemma is easy enough to prove. If B_1 is not contained in S then there is a $y \in B_1 - S$. Because $B_1 - S$ is open we can find an $r > 0$ such that the open ball B of radius r about y is contained in $B_1 - S$. We can then choose an $r_2 > 0$ which is less than r. Letting B_2 be the closed ball of radius r_2 about y we then have $B_2 \subseteq B \subseteq B_1 - S$. That concludes the proof of the lemma.

If S_1 contains no open ball then there is a closed ball B_1 of radius at most 1 whose intersection with S_1 is empty. If also S_2 contains no open ball then there is then a closed ball B_2 of radius at most $1/2$ contained in B_1 whose intersection with $S_2 \cap B_1$ is empty, and therefore whose intersection with $S_1 \cup S_2$ is empty. Continuing in this way, if no S_j contains an open ball then we obtain a nested sequence $B_1 \supseteq B_2 \supseteq \cdots$, of closed balls such that

$$B_n \cap \left(\cup_{j=1}^n S_j\right) = \emptyset$$

and

$$\delta(B_n) < 1/2^n.$$

By the lemma there is an

$$x \in \cap_{j\in\mathbf{Z}^+} B_j.$$

But then

$$x \notin \cup_{j=1}^n S_j$$

for all $n \in \mathbf{Z}^+$. Equivalently

$$x \notin \cup_{j\in\mathbf{Z}^+} S_j.$$

But

$$\cup_{j\in\mathbf{Z}^+} S_j = X,$$

so we have a contradiction, and our assumption that no S_j contains an open ball must be false. That concludes the proof of the theorem.

As an application of the Baire Category Theorem we show that there is a continuous nowhere differentiable function. Let $X_n \subseteq C([a, b])$, with $a < b$, consist of those functions x such that there is an $s \in [a, b]$ such that for all $t \in [a, b]$

$$|x(s) - x(t)| \leq n|s - t|.$$

Then X_n is closed. To see this, suppose $x_j \to y$ in $C([a, b])$ and $x_j \in X_n$ for each j. Because $x_j \in X_n$ there is an $s_j \in [a, b]$ such that for all $t \in [a, b]$ we have

$$|x_j(s_j) - x_j(t)| \leq n|s_j - t|.$$

The sequence $(s_1, s_2, \ldots)$ has a convergent subsequence $(s_{j_1}, s_{j_2} \ldots)$. Let

$$r = \lim_{k\to\infty} s_{j_k}.$$

Then, for all $k \in \mathbf{Z}^+$ and $t \in [a, b]$

$$|x_{j_k}(s_{j_k}) - x_{j_k}(t)| \leq n|s_{j_k} - t|.$$

Taking limits,

$$|y(r) - y(t)| \leq n|r - t|$$

for all $t \in [a, b]$ and hence

$$y \in X_n,$$

as claimed.

Also, X_n does not contain any open balls. To see this we make use of the 'sawtooth function'

$$w(t) = \begin{cases} 1 - 4k + 2t & \text{if } 2k - 1 \leq t \leq 2k, \\ 1 + 4k - 2t & \text{if } 2k \leq t \leq 2k + 1. \end{cases}$$

This function has the property that $\|w\| = 1$ and, for each $s \in \mathbf{R}$

$$\left| \frac{w(s) - w(t)}{s - t} \right| = 2.$$

for all t in a neighborhood of s. If $x \in C([a, b])$ then for every $\rho > 0$ there is, by the Weierstrass Approximation Theorem, a polynomial p whose restriction to $[a, b]$, which we will also call p, satisfies

$$\|x - p\| < \frac{\rho}{2}.$$

Choose

$$\gamma > \max_{[a,b]} |p'|$$

and define $y \in C([a, b])$ by

$$y(t) = p(t) + \rho 2 w(\kappa t),$$

where

$$\kappa = \frac{\gamma + n}{\rho},$$

and note that

$$\|x - y\| < \rho.$$

If $s, t \in [a, b]$ and $s \neq t$ then

$$|p(s) - p(t)| < \gamma|s - t|.$$

For any $s \in [a, b]$ there is a $t \in [a, b]$ such that $s \neq t$ and

$$|w(\kappa s) - w(\kappa t)| = 2\kappa|s - t|$$

and hence

$$|y(s) - y(t)| > (\kappa\rho - \gamma)|s - t| = n|s - t|.$$

Thus $y \notin X_n$. So the open ball of radius ρ about x is not contained in X_n. But x and ρ we arbitrary, so X_n contains no open balls.

By the Baire Category Theorem,

$$C([a, b]) \neq \cup_{n \in \mathbf{Z}^+} X_n.$$

But if x is differentiable even at a single point then

$$x \in \cup_{n \in \mathbf{Z}^+} X_n.$$

Indeed, if x is differentiable at $s \in [a, b]$ then the function

$$z = (t) = \begin{cases} \frac{x(s) - x(t)}{s - t} & \text{if } t \neq s, \\ x'(s) & \text{if } t = s \end{cases}$$

is continuous and hence bounded, so there is an n such that

$$|z(t)| \leq n$$

for all $t \in [a, b]$. But then $x \in X_n$. Thus the functions which are differentiable at at least one point form a subset of $\cup_{n \in \mathbf{Z}^+} X_n$, which is a *proper* subset of $C([a, b])$.

5.3 The Open Mapping and Closed Graph Theorems

Suppose that $A \in \mathcal{L}(E, F)$, where E and F are Banach spaces, is surjective. Let K_ρ and L_ρ be the open balls of radius ρ about 0 in E and F, respectively:

$$K_\rho = \{x \in E \colon \|x\| < \rho\}, \quad L_\rho = \{x \in E \colon \|x\| < \rho\}.$$

A consequence of the Baire Category Theorem is that there is an $\alpha > 0$ such that for all $\rho > 0$

$$L_{\alpha\rho} \subseteq \overline{A(K_\rho)}.$$

To prove this we observe that the surjectivity of A means that for all $y \in F$ there is an $x \in E$ with $Ax = y$. If $n > \|x\|$ then

$$y \in A(K_n) \subseteq \overline{A(K_n)}.$$

Since y was arbitrary,

$$F = \cup_{n\in\mathbf{Z}^+} \overline{A(K_n)}.$$

By the Baire Category Theorem there is at least one n for which $\overline{A(K_n)}$ contains an open ball. Let σ and z be the radius and centre of this ball. A is surjective, do $z = Aw$ for some $w \in E$. Then

$$L_\sigma \subseteq \overline{A(K_{n+\|w\|})}.$$

Let

$$\alpha = \frac{\sigma}{n + \|w\|}.$$

Then

$$L_{\alpha\rho} \subseteq \overline{A(K_\rho)}$$

for all $\rho > 0$.

The fact that we have the closure of $A(K_\rho)$ instead of $A(K_\rho)$ itself is a nuisance. We can get rid of it at the expense of replacing α by a smaller number β. More precisely, if $0 < \beta < \alpha$ then

$$L_{\beta\rho} \subseteq A(K_\rho).$$

To show this, let

$$\gamma = \frac{\alpha - \beta}{\alpha} < 1.$$

If

$$y \in L_{\beta\rho}$$

then

$$y \in \overline{A(K_{\beta\rho/\alpha})}.$$

In other words there is, for any $\epsilon_1 > 0$ a $w_1 \in K_{\beta\rho/\alpha}$ such that

$$\|y - Aw_1\| < \epsilon_1.$$

This holds in particular for

$$\epsilon_1 = \gamma\beta\rho.$$

So

$$y - Aw_1 \in L_{\gamma\beta\rho}.$$

But then

$$y - Aw_1 \in \overline{A(K_{\gamma\beta\rho/\alpha})}$$

and there is, for any $\epsilon_2 > 0$ a $w_2 \in K_{\gamma\beta\rho/\alpha}$ such that

$$\|y - Aw_1 - Aw_2\| < \epsilon_2.$$

This holds in particular for

$$\epsilon_2 = \gamma^2\beta\rho.$$

So

$$y - Aw_1 - Aw_2 \in L_{\gamma^2\beta\rho}.$$

We continue on in this way, obtaining a sequence $(w_1, w_2, \ldots)$ in E such that

$$\|w_n\| < \gamma^{n-1}\beta\rho/\alpha$$

and

$$\|y - A\sum_{j=1}^{n} w_j\| < \gamma^n\beta\rho.$$

By the comparison test, the sum of the w's converges. Let

$$x = \sum_{j=1}^{\infty} w.$$

Then

$$\|x\| < \sum_{j=1}^{\infty} \gamma^{n-1}\beta\rho/\alpha = \rho$$

and

$$y = Ax$$

In other words,

$$y \in A(K_\rho).$$

But $y \in L_{\beta\rho}$ was arbitrary, so

$$L_{\beta\rho} \subseteq A(K_\rho).$$

Suppose that $U \subset E$ is open and $y \in A(U)$. Then $y = Ax$ for some $x \in U$ and, because U is open, there is a $\delta > 0$ such that the ball of radius δ about x is contained in U. Then the ball of radius δ/β is contained in $A(U)$, so every $y \in A(U)$ is contained in an open ball in $A(U)$. Mappings, i.e. functions, between topological spaces with the property that the image of every open set is open are called *open mappings*. What we've just shown is that every bounded linear transformation between Banach spaces which is surjective is an open mapping. This is called the Open Mapping Theorem.

Suppose that E and F are Banach spaces and $A \in \mathcal{L}(E, F)$ is bijective. Then A, considered as a mapping between sets, has an inverse, which we will call A^{-1}. It's straightforward to check that A^{-1} is a linear transformation, but it's not immediately obvious that this inverse is continuous. Indeed there are continuous mappings from metric space to a metric space which are bijective, but whose inverses are not continuous. But fortunately this does not happen for A. If $U \in E$ is open and $V \in F$ is the pre-image of U under A^{-1} then $V = A(U)$ and this, by the Open Mapping Theorem, is an open set.

If $f: E \to F$ is a function then its graph is, by definition, the set

$$G = \{(x, y) \in E \times F : y = f(x)\}.$$

We can make $E \times F$ into a Banach space by means of the norm

$$\|(x, y)\|_{E\times F} = \|x\|_E + \|y\|_F.$$

If f is continuous then G is closed. To see this, suppose $((x_1, y_1), (x_2, y_2), \ldots)$ is a sequence in $E \times F$ converging to (w, z), with (x_j, y_j) in G for each j. Then

$$\lim_{n\to\infty} \|(x_n, y_n) - (w, z)\| = 0$$

or

$$\lim_{n\to\infty} (\|x_n - w\| + \|y_n - z\|) = 0$$

and hence

$$\lim_{n\to\infty} \|x_n - w\| = 0, \quad \lim_{n\to\infty} \|y_n - z\| = 0$$

and

$$\lim_{n\to\infty} x_n = w, \quad \lim_{n\to\infty} y_n = z.$$

Since $(x_n, y_n) \in G$ we have $y_n = f(x_n)$ and the last of the equations above can be written as

$$\lim_{n\to\infty} f(x_n) = z.$$

By the continuity of f we then have

$$f(w) = f\left(\lim_{n\to\infty} x_n\right) = \lim_{n\to\infty} f(x_n) = z.$$

So

$$(w, z) \in G.$$

The limit of any sequence in G is therefore also in G, so G is closed, as claimed.

For linear transformations between Banach spaces the converse holds: If $A: E \to F$ is linear and G is closed then A is continuous. To see this, let $P_E: G \to E$ and $P_F: G \to F$ be the projections

$$P_E((x, y)) = x, \quad P_F((x, y)) = y.$$

These are linear and bounded, with $\|P_E\| = \|P_F\| \leq 1$. Also P_E is bijective, by the definition of a function! So by the corollary to the Open Mapping Theorem P_E has continuous inverse P_E^{-1}. But

$$A = P_F P_E^{-1},$$

so A is continuous.

We've just seen that a linear transformations from a Banach Space to a Banach space is continuous if and only if its graph is closed. This statement is called the Closed Graph Theorem.

5.4 Banach-Steinhaus

The Uniform Boundedness Principle says the if E and F are Banach spaces and $S \subseteq \mathcal{L}(E, F)$ is such that for all $x \in E$

$$\sup_{A\in S} \|Ax\| < \infty$$

then

$$\sup_{A\in S} \|A\| < \infty.$$

Let E_n be the subset of E consisting of those x for which

$$\sup_{A\in S} \|Ax\| \leq n.$$

Now

$$E_n = \cap_{A\in S}\{x \in E : \|Ax\| \leq n\}$$

is an intersection of closed sets and hence is closed. Also

$$\cup_{n\in\mathbf{Z}^+} E_n = E.$$

By the Baire Category Theorem there is an open ball $B \subseteq E_n$ for some n. Let w be the centre of B and let ρ be its radius. If $\|x\| \leq 1$ and $A \in S$ then

$$w \in B \subseteq E_n, \quad w + \frac{\rho}{2} x \in B \subseteq E_n$$

and

$$\begin{aligned}\|Ax\| &= \left\| \frac{2}{\rho}\left(A\left(w+\frac{\rho}{2}x\right) - Aw\right)\right\| \\ &\le \frac{2}{\rho}\left(\left\|A\left(w+\frac{\rho}{2}x\right)\right\| + \|Aw\|\right) \\ &\le \frac{2}{\rho}(n+n) = \frac{4n}{\rho}.\end{aligned}$$

This holds for all x with $\|x\| \le 1$, so

$$\|A\| \le \frac{4n}{\rho}.$$

This holds for all $A \in S$, so

$$\sup_{A\in S} \|A\| \le \frac{4n}{\rho} < \infty.$$

This completes the proof of the theorem.

Suppose that $(A_1, A_2, \ldots)$ is a sequence in $\mathcal{L}(E, F)$ and

$$\lim_{n\to\infty} A_n x$$

exists for each $x \in E$. Define $B\colon E \to F$ by

$$Bx = \lim_{n\to\infty} A_n x.$$

Then, by the properties of limits, B is linear. But is it bounded? In general the answer is no, but if E is a Banach space then the Uniform Boundedness Principle shows that the answer is yes. Convergent sequences are bounded, so

$$\sup_{n\in\mathbf{Z}^+} \|A_n x\| < \infty$$

for each $x \in E$. By the Uniform Boundedness Principle,

$$\sup_{n\in\mathbf{Z}^+} \|A_n\| < \infty$$

so there is a $\gamma \in \mathbf{R}^+$ such that

$$\|A_n\| \le \gamma$$

for all $n \in \mathbf{Z}^+$. In other words, pointwise convergence implies boundedness. Now

$$\|A_n x\| \le \|A_n\|\|x\| \le \gamma\|x\|$$

for all $n \in \mathbf{Z}^+$ and $x \in E$. Taking limits,

$$\|Bx\| \le \gamma\|x\|$$

for all $x \in E$. Thus γ is a bound for B, and so B is bounded. In fact we can take $\gamma = \sup \|A_n\|$, so

$$\|B\| \le \sup \|A_n\|.$$

With a little more effort we can replace the sup with a lim sup.

We saw in the course of the proof that pointwise convergence of a sequence in $\mathcal{L}(E, F)$ implies boundedness. Equivalently, unboundedness implies a failure of pointwise convergence. We can use this to show that there are periodic continuous functions whose Fourier series fail to converge somewhere. Suppose x is periodic with period 2. Its Fourier coefficients are then

$$\alpha_m = \frac{1}{2}\int_{-1}^{1} x(t)\exp(-\pi i m t)\,dt$$

and its n'th Fourier approximand is

$$\sum_{m=-n}^{n} \alpha_m \exp(\pi i m t) = \int_{-1}^{1} k_n(s-t)x(t)\,dt$$

where

$$k_n(t) = \frac{\sin\left(\left(n+\frac{1}{2}\right)\pi t\right)}{2\sin\left(\frac{1}{2}\pi t\right)}.$$

Note that k_n is even. We will show that there is a continuous periodic function x of period 2 whose Fourier approximands fail to converge to x at $s = 0$. Let E be the subspace of $C([-1, 1])$ consisting of functions for which $x(-1) = x(1)$, i.e. of functions which can be extended to periodic continuous function of period 2, and let $F = \mathbf{C}$. Let $A_n\colon E \to F$ defined by

$$A_n x = \int_{-1}^{1} k_n(t)x(t)\,dt$$

Then

$$\begin{aligned}\|A_n\| &= \int_{-1}^{1} |k_n(t)|\,dt \\ &\ge \int_{-1}^{1} \frac{\left|\sin\left(\left(n+\frac{1}{2}\right)\pi t\right)\right|}{2\left|\frac{1}{2}\pi t\right|}\,dt \\ &= \frac{1}{\pi}\int_{-(n-1/2)\pi}^{(n+1/2)\pi} \left|\frac{\sin s}{s}\right|\,ds \to \infty.\end{aligned}$$

If $(A_1x, A_2x, \ldots)$ converged for each x then the sequence $(\|A_1\|, \|A_2\|, \ldots)$ would have to be bounded, so there is at least one continuous function on $[-1,1]$ whose Fourier series is not convergent at 0.

Recall again that pointwise convergence of a sequence in $\mathcal{L}(E,F)$ implies boundedness. The converse is false, but the Banach-Steinhaus Theorem provides a sort of converse, with the additional assumption of pointwise convergence on a dense subspace. More precisely, suppose E and F are Banach spaces, D is a subspace of E, $\overline{D} = E$, and $(A_1, A_2, \ldots)$ is a bounded sequence in $\mathcal{L}(E,F)$ such that

$$\lim_{n\to\infty} A_n x$$

exists for all $x \in D$. Then there is a $B \in \mathcal{L}(E,F)$ such that

$$\lim_{n\to\infty} A_n x = Bx$$

for all $x \in E$.

To prove the Banach-Steinhaus Theorem we consider an arbitrary $x \in E$ and $\epsilon > 0$. Let

$$\gamma = \sup \|A_n\|$$

and choose $y \in D$ such that

$$\|x - y\| < \frac{\epsilon}{3\gamma}$$

Because $y \in D$ the sequence $(A_1y, A_2y, \ldots)$ is convergent and hence Cauchy. Choose N such that for all $j, k > N$ we have

$$\|A_j y - A_k y\| < \frac{\epsilon}{3}.$$

Then

$$\begin{aligned}
&\|A_j x - A_k x\| \\
&\quad = \|A_j(x-y) + (A_j y - A_k y) - A_k(x-y)\| \\
&\quad \le \|A_j(x-y)\| + \|A_j y - A_k y\| + \|A_k(x-y)\| \\
&\quad \le \|A_j\|\|x-y\| + \|A_j y - A_k y\| + \|A_k\|\|x-y\| \\
&\quad < \gamma\frac{\epsilon}{3\gamma} + \frac{\epsilon}{3} + \gamma\frac{\epsilon}{3\gamma} = \epsilon,
\end{aligned}$$

so $(A_1x, A_2x, \ldots)$ is Cauchy and hence, since we assumed F was a Banach space, convergent. This holds for all $x \in E$, so the preceding corollary to the Uniform Boundedness Principle shows that there is a $B \in \mathcal{L}(E,F)$ such that

$$\lim_{n\to\infty} A_n x = Bx$$

for all $x \in E$.

5.5 Krein-Milman

If K is a convex subset of a vector space then it contains the midpoint of any two distinct in K. It may happen that all points in K are midpoints of distinct points in K. In any non-trivial Banach space the unit ball has this property, for example. Any point which is not the midpoint of a pair of distinct elements of K is called an *extreme point*. As we've just seen, there need not be any extreme points.

If S is a subset of a normed space E then we define the *closed convex hull* of S to be the intersection of all closed convex subsets of E which contain S. The closed convex hull is itself closed and convex, because the intersection of closed sets is closed and the intersection of convex sets is convex. The closed convex hull of S is in fact the smallest closed convex subset of E containing S, because the intersection of a collection of sets is a subset of any element of the collection.

The Krein-Milman Theorem states that if K is a compact convex subset of a Banach space E then K is the closed convex hull of its set of extreme points. There is no loss of generality in assuming that $\mathbf{K} = \mathbf{R}$, so we will assume that in the proof.

Our proof of the Krein-Milman theorem will rely on the Hahn-Banach Separation Theorem, which says that if E is a real Banach space, U is a non-empty open convex subset of E, V is a non-empty convex subset of E and $U \cap V = \emptyset$ then there is a $g \in E'$ and $\gamma \in \mathbf{R}$ such that for all $x \in U$ and $y \in V$ we have $g(x) > \gamma \ge g(y)$. The proof is based on the seminorm version of the Hahn-Banach Theorem. We choose $u \in U$ and $v \in V$ and define S to be the set of points in E of the form

$$z = x - y - u + v$$

with $x \in U$ and $v \in V$. We then define

$$p(z) = \inf\{t \in \mathbf{R}: t \geq 0, z \in tS\}$$

and show that p is a semi-norm on E. If F is the span of the single vector $u - v$ and $f \in F'$ is the unique function with

$$f(u - v) = 1$$

then

$$|f(z)| \leq p(z)$$

for all $x \in F$. The Hahn-Banach gives us a $g \in E'$ such that

$$g(z) = f(z)$$

for $z \in F$ and

$$|g(z)| \leq p(z)$$

for $z \in E$. Then

$$g(x) > g(y)$$

for all $x \in U$ and $y \in V$ and so, taking

$$\gamma = \inf_{x \in U} g(x),$$

we get

$$g(x) > \gamma \geq g(y)$$

for all $x \in U$ and $y \in V$. This completes the proof of the Hahn-Banach Separation Theorem.

The theorem has a useful corollary. Suppose V is a non-empty closed convex subset of a Banach space E and $z \notin V$. Then there is an open ball U with centre z and radius ρ such that $U \cap V = \emptyset$. Applying the theorem we have $g \in E'$ and $\gamma \in \mathbf{R}$ such that

$$g(x) > \gamma \geq g(y)$$

for all $x \in U$ and $y \in V$. Then

$$\inf_{x \in U} g(x) \geq \gamma.$$

Now

$$\|g\| = \sup_{\|u\| < 1} g(u) = \sup_{\|z - x\| < \rho} \frac{g(z) - g(x)}{\rho}$$

and hence

$$\inf_{x \in U} g(x) = g(z) - \|g\|\rho.$$

If

$$\beta = \gamma - \frac{1}{2}\|g\|\rho$$

then

$$g(z) > \beta > g(y)$$

for all $y \in V$.

If the Krein-Milman Theorem is correct then it follows in particular that K has extreme points, which is not obvious. We don't however need the Krein-Milman Theorem to prove this. We say that a closed non-empty subset of $S \subset K$ is extreme if

$$s, t \in \mathbf{R}^+, \quad x, y \in K, \quad \frac{sx + ty}{s + t} \in S$$

implies

$$x, y \in S.$$

Note that K itself is extreme. The terminology is somewhat misleading because not every point in an extreme set is necessarily extreme. We'll soon see though that at least one of them is. We can order the extreme sets by reverse inclusion. The intersection of a totally ordered collection of extreme sets is extreme. The compactness of K ensures that the intersection is non-empty. We can then use Zorn's Lemma to show that there is a maximal extreme set. 'Maximal' here means that no proper subset is extreme. This maximal extreme set must be a single point. To see this, suppose M is maximal and $x, y \in M$. If $x \neq y$ then, by the corollary to the Hahn-Banach Separation Theorem, there is a $g \in E'$ such that $g(x) > g(y)$. The set

$$\{z \in M: g(z) \geq g(x)\}$$

would be a smaller extreme set, contradicting maximality. So we have at least one maximal extreme set and this set contains only exactly one point. That point must be an extreme point. In fact we have something a bit stronger. We could apply Zorn's Lemma not to all extreme sets but to all extreme sets contained in any given one, so every extreme set contains an extreme point, as mentioned above.

The set of extreme points of K is contained in K. Since K is closed and convex the closed convex hull of the set of extreme points is also contained in K. The non-trivial part of Krein-Milman is the reverse inclusion. Let V be the closed convex hull of the extreme points. We need to show that $K \subseteq V$.

There is at least one extreme point, so V is a non-empty closed convex set. If $z \in K - V$ then there is, by the corollary to the Hahn-Banach Separation Theorem, a $g \in E'$ and $\beta \in \mathbf{R}$ such that

$$g(z) > \beta > \max_V g.$$

But then

$$\{x \in K : g(x) = \max_K g\}$$

would be an extreme set without extreme points, which we've seen is impossible. So there is no $z \in K - V$ and therefore $K \subseteq V$. That concludes the proof of the Krein-Milman Theorem.

As an application, we consider a class of combinatorial optimisation problems called assignment problems. The classical version of this is assigning people to tasks. We have n people and n tasks. The productivity of the j'th person performing the k'th task is $p_{j,k}$. Let $x_{j,k}$ be 1 if the j'th person is assigned to the k'th task and 0 otherwise. We want each person assigned to one and only one task,

$$\sum_{j=1}^{n} x_{j,k} = 1$$

and we want each task to have one and only person assigned to it,

$$\sum_{k=1}^{n} x_{j,k} = 1.$$

Among all such assignments we want to select one which maximises the total productivity

$$\sum_{j=1}^{n} \sum_{k=1}^{n} p_{j,k} x_{j,k}.$$

Possible assignments correspond to permutations of $\{1, 2, \ldots, n\}$, so we could just check all the possibilities, but there are $n!$ of them and that's not practical for large n.

Linear programming problems are ones where we seek to maximise a linear function of finitely many real variables subject to finitely many inequality constraints on those variables. The assignment problem looks like a linear programming problem, but the constraints are equations and the variables are restricted to the set $\{0, 1\}$.

A first step in converting our assignment problem to a linear programming problem is to rewrite the constraints as

$$x_{j,k} \geq 0, \quad \sum_{j=1}^{n} x_{j,k} \geq 1, \quad \sum_{k=1}^{n} x_{j,k} \leq 1.$$

Any possible assignment satisfies these inequalities. Conversely, suppose the inequalities are satisfied. Summing

$$\sum_{j=1}^{n} x_{j,k} \geq 1$$

over k gives

$$\sum_{j=1}^{n} \sum_{k=1}^{n} x_{j,k} \geq n,$$

while summing

$$\sum_{k=1}^{n} x_{j,k} \leq 1$$

over j gives the reverse inequality. The only way to satisfy both is if all the inequalities involved are actually equations, so

$$x_{j,k} \geq 0, \quad \sum_{j=1}^{n} x_{j,k} = 1, \quad \sum_{k=1}^{n} x_{j,k} = 1.$$

Also,

$$x_{j,k} \leq \sum_{l=1}^{n} x_{j,l} = 1,$$

since $x_{j,l} \geq 0$ for $l \neq k$. So if $x_{j,k} \in \mathbf{Z}$ for each j, k then

$$x_{j,k} \in \{0, 1\}, \quad \sum_{j=1}^{n} x_{j,k} = 1, \quad \sum_{k=1}^{n} x_{j,k} = 1.$$

In other words, it represents a possible assignment of people to tasks. There are, however, plenty of non-integer solutions to the inequalities, e.g. $x_{j,k} = 1/n$.

The set

$$K = \left\{ x \in \mathbf{R}^{n^2} : x_{j,k} \geq 0, \sum_{j=1}^{n} x_{j,k} \geq 1, \sum_{k=1}^{n} x_{j,k} \leq 1 \right\}$$

is compact and the function

$$f(x) = \sum_{j=1}^{n} \sum_{k=1}^{n} p_{j,k} x_{j,k}$$

is continuous, so it has a maximum. The set

$$K_f = \left\{ x \in K \colon f(x) = \max_K f \right\}$$

is compact and convex, so by the Krein-Milman Theorem it is the convex hull of its extreme points. In particular, it has an extreme point. Any extreme point has $x_{j,k} \in \{0,1\}$ though.

First of all, if x is an extreme point of K_f then it must also be an extreme point of K. This isn't simply a consequence of $K_f \subseteq K$, but it is nonetheless true. If x is an extreme point of K_f but not of K then

$$x = \frac{y+z}{2}, \quad y, z \in K$$

with $y \neq z$. Then

$$f(x) = \frac{1}{2} f(y) + \frac{1}{2} f(z).$$

So either

$$f(x) = f(y) = f(z),$$

in which case

$$y, z \in K_f,$$

which is impossible because x is an extreme point of K_f, or one of $f(y) > f(x)$ or $f(z) > f(x)$ is true. But

$$f(x) = \max_K f,$$

so this too is impossible. So we only need to consider extreme points of K. Given $x \in K$ we construct two permutations $\sigma, \tau \in S_n$ as follows. Let

$$J_1 = K_1 = \{1, 2, \ldots n\}$$

Choose $j_1 \in J_1$ and $k_1 \in K_1$ to make

$$x_{j_1,k_1} = \max j \in J_1, k \in K_1 x_{j,k}$$

There may be multiple choices which work, but just choose one and Set

$$J_2 = J_1 - j_1, \quad K_2 = K_1 - k_1$$

and

$$\sigma(1) = j_1, \quad \tau(1) = k_1.$$

Choose $2_1 \in J_2$ and $k_2 \in K_2$ to make

$$x_{j_2,k_2} = \max j \in J_2, k \in K_2 x_{j,k}$$

Set

$$J_3 = J_2 - j_2, \quad K_3 = K_2 - k_2$$

and

$$\sigma(2) = j_2, \quad \tau(2) = k_2.$$

Continue in this way until we've chosen $\sigma(n)$ and τ_n. The effect of this is find a permutation σ of rows and a permutation τ of columns such that the permuted version of x has the property that each of the submatrices

$$\begin{pmatrix} x_{\sigma(m),\tau(m)} & \cdots & x_{\sigma(m),\tau(n)} \\ \vdots & \ddots & \vdots \\ x_{\sigma(n),\tau(m)} & \cdots & x_{\sigma(n),\tau(n)} \end{pmatrix}$$

has its largest entry in its upper left hand corner. If

$$x_{\sigma(n),\tau(n)} < 1$$

then set

$$\delta = \max\left(1, \frac{x_{\sigma(n),\tau(n)}}{1 - x_{\sigma(n),\tau(n)}}\right),$$

$$y_{j,k} = \begin{cases} (1+\delta) x_{j,k} - \delta & \text{if } j = k, \\ (1+\delta) x_{j,k} & \text{if } j \neq k, \end{cases}$$

$$z_{j,k} = \begin{cases} (1-\delta) x_{j,k} + \delta & \text{if } j = k, \\ (1-\delta) x_{j,k} & \text{if } j \neq k. \end{cases}$$

Then

$$y, z \in K, \quad x = \frac{y+z}{2},$$

so x is not an extreme point of K. So any extreme point x has

$$x_{\sigma(n),\tau(n)} = 1.$$

But then

$$x_{\sigma(j),\tau(k)} = \begin{cases} 1 & \text{if } j = k, \\ 0 & \text{if } j \neq k \end{cases}$$

for all $j \in \{1, \ldots, 1\}$ and x is the permutation matrix

$$x_{l,m} = \begin{cases} 1 & \text{if } \sigma^{-1}(l) = \tau^{-1}(m), \\ 0 & \text{if } \sigma^{-1}(l) \neq \tau^{-1}(m). \end{cases}$$

We've just seen that there is at least point in K_f, i.e. one solution of the linear programming problem, which is an extreme point of K and that this must be given by an element of S_n, i.e. a solution of the assignment problem.

It's not necessarily true that every solution of the linear programming problem is a solution of the assignement problem, but there are efficient methods of solving the linear programming problem which are guaranteed to find such a solution. Alternatively,[9] we can first find a w which maximises

$$\sum_{j=1}^{n} \sum_{k=1}^{n} p_{j,k} w_{j,k}$$

$$w_{j,k} \geq 0, \quad \sum_{j=1}^{n} w_{j,k} \geq 1, \quad \sum_{k=1}^{n} w_{j,k} \leq 1$$

and then, if w itself is not a solution to the assignment problem, find an x which maximises

$$\sum_{j=1}^{n} j e^k k x_{j,k}$$

subject to

$$x_{j,k} \geq 0, \quad \sum_{j=1}^{n} x_{j,k} \geq 1, \quad \sum_{k=1}^{n} x_{j,k} \leq 1,$$

$$\sum_{j=1}^{n} \sum_{k=1}^{n} p_{j,k} x_{j,k} \geq \sum_{j=1}^{n} \sum_{k=1}^{n} p_{j,k} w_{j,k}.$$

We know the constraints in the second problem are satisfiable because x satisfies them. Anything which satisfies the second problem is a solution to the first problem as well, and the solution to the second problem must occur at an extreme point because the objective function takes a different value at each extreme point.

[9] The method which follows works with exact arithmetic, but needs some modification to work with floating point arithmetic.