MATH 442

Analysis II: Topics in Analysis

2008

§0. Introduction.

As it says in the Prospectus, there are several possible topics that might be treated in this course, and I cannot pretend to have made an "authoritative" selection. But I hope that all the matters I shall mention are interesting in themselves; and they certainly have deep and striking applications, mostly beyond the scope of our course.

The course is an introduction to functional analysis, and it may be appropriate to say briefly what that phrase connotes. (As with many other mathematical topics, it is less precise now than it once was.)

In "classical" analysis, one studies individual functions; an example of a typical theorem might be, let us say, Lagrange's form of Taylor's theorem, which states a property that a suitable function of a real variable has. But, about a century ago, it became clear that there were some problems — I think they were mostly related to partial differential equations, for instance to solutions of the Cauchy-Riemann equations — in which it was helpful to consider functions of particular types in the mass: that is, to consider a class of functions as a "space" in its own right, perhaps a metric space or a vector space or both. The idea had really been around for quite a long time without being very precisely formulated; at the beginning of last century the abstract concepts that were needed were at last introduced. Thus, my primary aim is to convince you that this is a valuable procedure, and to present some of the basic results.

Before beginning, one general apology is appropriate. I shall, for the most part, focus on metric and on Banach spaces, and I shall often give the most "basic" versions of results rather than the "optimal" formulations. There are important reasons why more general spaces or more general results are used, but it is difficult to grasp the ideas behind the generalizations until you are comfortable with simpler cases in which the proofs are more natural. But it is not clear to me precisely what we may assume; so you may consider I am very casual in some respects and absurdly pernickety in others, and there will be some duplication of things I have taught elsewhere. Do not hesitate to complain when it seems appropriate.

There are many textbooks which could be mentioned. The subject still has a certain veneer of modernity that makes it attractive to authors. I shall not follow any book in particular, and not all the topics I shall mention may appear in any given book. An excellent reference is Kolmogorov and Fomin (which also gives a very readable account of some topics I shall more or less take for granted). Others are Simmonds, Taylor, Yosida (long, but also very compressed because it discusses many important applications), Naimark (Normed Rings; especially chapter 1), Dunford and Schwartz (very, very long; but especially parts of vol. 1).

To motivate our work in a general way — for I shall never deal with the problems that follow in any detail — let me briefly describe a couple of problems of historical interest.

Suppose we have two points A and B in space. To avoid trivializing the question, A is to be higher than B and not in the same vertical line; we assume the gravitational field is entirely uniform for mathematical simplicity. We want to join A and B by a smooth rail, so that we can slide or roll down the rail from A to B, starting from rest at A, under the influence of gravity alone and with negligible friction. It is clear that the curve of the rail from A to B does affect the time it takes to reach B: if the rail slopes upwards at the start, our journey never begins (and if the rail ever returns to the height of A, the motion stops); or if, at the beginning, the rail follows a very shallow downwards spiral, a rather long distance might have to be traversed at a very low speed; or the rail might go steeply downwards at the start, so that the journey may be taken at a rather high speed. Which curve from A to B would give the shortest possible time for the journey? (This is the *brachistochrone problem*, which was current in the 17th century and was solved by several of its greatest names at the very end of the century. Brachistos is "shortest" and chronos is "time".)

This is, obviously, a question of the "find a maximum or minimum" type. But in this case the number we want to minimize, the travel time, depends, not on one or several real variables, but on a whole curve in three-dimensional space. If we describe the possible curves parametrically, say as functions $p : [0,1] \longrightarrow \mathbb{R}^3$, where $p(0) = A$ and $p(1) = B$, then the time t that the journey takes is a function for which the independent variable is the function p determining the curve, and not any finite n-tuple of real numbers. These functions p, or the curves they describe, clearly cannot be specified by any fixed finite number of numerical parameters, and in this vague sense the domain of the function t is an "infinite-dimensional space" whose "points" are the *functions* p.

Another even more familiar problem is this. Suppose that we have a certain area of material which we can use to form a surface in space (for instance a certain volume of plastic, which must be spread to a definite very small thickness to construct the surface; what I have in mind is that the shape of the surface may be freely varied, but its total area cannot). What is the largest volume that can be enclosed within the surface? Of course we all know the answer — the material must form a sphere if the enclosed volume is to be the largest possible. But it is not easy to explain why this is the answer. Once again, the volume V enclosed is a function whose argument is the "surface"; the domain of V is the set of possible surfaces, which is an "infinite-dimensional space".

In both these examples, there is a question that arises before the main problem. Have we any reason to suppose that there *exists* a best possible curve in the brachistochrone problem, or a best possible surface in the volume problem? It is at least conceivable that there is no "briefest possible time" for rolling from A to B; there might be a whole sequence of curves, giving us successively shorter and shorter travel times approaching a limit, but no curve that would exactly realize that limit. One might imagine the sequence of curves that give shorter and shorter travel times wiggling about in space and not approaching any sort of limiting curve. This is a serious matter; the assumption that the theoretical optimum must be "attained", that is, be a genuine value given by some curve or surface or whatever, is not always justified.

The general moral is that there is good reason to consider "infinite-dimensional" objects. They are not artificial constructs; despite the strange images the phrase "infinite-dimensional" may suggest, it really denotes no more than the impossibility of describing the things we are considering by a finite number of parameters. Our aim, in fact, is to develop theories that will make it possible to argue with such objects very much as we argue with $\mathbb{R}^n$ or $\mathbb{C}^n$. To do this, we must first decide what properties of $\mathbb{R}^n$ or $\mathbb{C}^n$ we regard as really significant, and we must necessarily do this in rather abstract terms, precisely because the objects that interest us are difficult to describe directly in the elementary terms we are used to.

§1. Metric spaces and mappings

Definition 1.1. Let Ω be any set. A *metric* (or *distance function*) on Ω is a function $d : \Omega \times \Omega \longrightarrow \mathbb{R}$ such that, for any $x, y, z \in \Omega$,

(a) $d(x, y) = 0$ if and only if $x = y$, and
(b) $d(x, z) \leq d(x, y) + d(z, y)$.

(a) and *(b)* imply the further properties of the metric

(c) $(\forall x, y \in \Omega) \quad d(x, y) \geq 0$,
(d) $(\forall x, y \in \Omega) \quad d(x, y) = d(y, x)$,

which are often taken as part of the definition; and *(b)* may then be written as

(e) $d(x, z) \leq d(x, y) + d(y, z)$.

The pair (Ω, d) is called a *metric space*. If d has been unambiguously fixed, one may speak of "the metric space Ω" and suppress mention of d. *(b)*, or *(e)*, is the *triangle inequality* for d.

The concept of a metric space arose from the examples of $\mathbb{R}^n$ or $\mathbb{C}^n$, with their "Euclidean" or "Hermitian" or "standard" metrics

$$d((p_1, p_2, \ldots, p_n), (q_1, q_2, \ldots, q_n)) := \sqrt{\sum_{k=1}^{n} |p_k - q_k|^2}$$

(for $\mathbb{R}^n$, the ps and qs are real numbers; for $\mathbb{C}^n$, they are complex numbers).

Definition 1.2. Let d be a metric on Ω, and suppose that Ψ is a subset of Ω. Define

$$d_\Psi : \Psi \times \Psi \longrightarrow \mathbb{R} : (x, y) \mapsto d(x, y)$$

(which makes sense, as x and y are points of Ω as well). Then d_Ψ is a metric, the subspace metric, in Ψ, and the metric space (Ψ, d_Ψ) is called a *metric subspace* of (Ω, d).

For obvious reasons, one usually writes d instead of d_Ψ, and says "Ψ is a (metric) subspace of the metric space Ω".

Definition 1.3. A sequence (x_n) in the metric space (Ω, d) is said to *converge* to $x \in \Omega$ (as $n \to \infty$) if

$$(\forall \epsilon > 0)(\exists N \in \mathbb{N}) \quad n \geq N \Longrightarrow d(x_n, x) < \epsilon .$$

One may say "x_n (or (x_n)) tends to x", and "x is the limit of x_n" (or of (x_n)), and one writes $x_n \to x$ (as $n \to \infty$).

Thus $x_n \to x$ if and only if the numerical sequence $(d(x_n, x))$ tends to 0. Notice that the "N" whose existence is demanded by the definition will probably have to be increased if ϵ is diminished, and that it will also depend on the specific sequence under consideration.

Lemma 1.4. *A sequence in* Ω *can have at most one limit.* $\square$

It may of course not converge to anything, but if it has a limit, that limit is unique.

Definition 1.5. Let (Ω, d_Ω) and (Ψ, d_Ψ) be metric spaces, let $f : \Omega \longrightarrow \Psi$ be a mapping, and suppose $x \in \Omega$. f is *continuous* at $x \in \Omega$ if

$$(\forall \epsilon > 0)(\exists \delta > 0) \quad y \in \Omega \ \& \ d_\Omega(x, y) < \delta \Longrightarrow d_\Psi(f(x), f(y)) < \epsilon \,. \tag{1}$$

We also say that x is a *point of continuity* of f. If $A \subseteq \Omega$ and every point of A is a point of continuity of f, we say f is continuous on A; if f is continuous on Ω, we say f is *continuous*, without qualification.

Again, the definition requires that for each positive real number ϵ there should exist a corresponding positive real number δ; if ϵ is decreased, then δ will probably have to be reduced too; and the choice of δ will also depend on the function f and the point x.

Lemma 1.6. *f is continuous at x if and only if, for any sequence (x_n) in Ω that converges to x, the sequence $f(x_n)$ in Ψ converges to $f(x)$.* $\square$

This lemma need not hold in general topological spaces.

Example 1.7. A simple example: take $\Omega := (0,1)$ as a metric subspace of $\mathbb{R}$, and $\psi := \mathbb{R}$, in both cases with the standard metric, and let $f(x) := 1/x$ for each $x \in (0,1)$. Given $x \in (0,1)$ and $\epsilon > 0$, we want to see if there can be a $\delta > 0$ such that

$$\text{whenever } y \in (0,1) \text{ and } |x - y| < \delta \text{, then } \left| \frac{1}{x} - \frac{1}{y} \right| < \epsilon \,; \tag{2}$$

that is, $|x - y| < \epsilon |xy|$. Clearly, then, we require: $\delta > 0$, $\delta < x$ (since otherwise y could be arbitrarily small and $1/y$ arbitrarily large), and $\delta \le \epsilon x(x - \delta)$, as we may arrange for y to be as close as we like to $x - \delta$. The last condition means that

$$\delta \le \frac{\epsilon x^2}{1 + \epsilon x} \,,$$

and it is easily checked that $\delta := \epsilon x^2/(1 + \epsilon x)$ does indeed satisfy (2). This is, therefore, the "best possible", i.e. the largest, choice of δ in this specific problem; it is clearly both positive and less than x. [In more complicated examples, it is usually impossible to calculate the largest value of δ explicitly.] But it is obvious that it depends, not only on ϵ and on the function considered, but also on x.

Definition 1.8. Let $\mathcal{F}$ be a family of mappings from Ω to Ψ. The family $\mathcal{F}$ is *equicontinuous at x* if

$$(\forall \epsilon > 0)(\exists \delta > 0)(\forall f \in \mathcal{F}) \quad y \in \Omega \ \& \ d_\Omega(x, y) < \delta \Longrightarrow d_\Psi(f(x), f(y)) < \epsilon \,. \tag{3}$$

Thus, equicontinuity of $\mathcal{F}$ at x means that, for any positive ϵ, there is a δ that "works" for that ϵ and the given point x *and for all of the functions in $\mathcal{F}$ simultaneously*. Such a δ will quite probably not be "optimal" for any of the functions considered. Evidently equicontinuity at x of a family that consists of a single function f, $\mathcal{F} = \{f\}$, is just the same as continuity of f at x; and any *finite* family of functions, $\mathcal{F} = \{f_1, f_2, \ldots, f_k\}$, each one of which is continuous at x, is equicontinuous at x. (For a given ϵ, take δ_i to satisfy (1) for f_i, and then let $\delta := \min\{\delta_1, \delta_2, \ldots, \delta_k\}$. Then δ satisfies (3).)

Example 1.9. It is easy to construct families of mappings that are *not* equicontinuous at a point; for instance, the family of functions $\{f_\alpha : \alpha > 0\}$, where

$$f_\alpha : (0,1) \longrightarrow \mathbb{R} : x \mapsto \alpha/x \tag{4}$$

for each α, is not equicontinuous at any point of $(0,1)$.

Definition 1.10. *(a)* Let A be a subset of Ω, and $f : \Omega \longrightarrow \Psi$. f is *uniformly continuous on A* if

$$(\forall \epsilon > 0)(\exists \delta > 0)(\forall a \in A) \quad y \in \Omega \ \& \ d_\Omega(a,y) < \delta \Longrightarrow d_\Psi(f(a), f(y)) < \epsilon .$$

If $A = \Omega$, we may say that f is uniformly continuous without qualification.

(b) Let $\mathcal{F}$ be a family of functions from Ω to Ψ. We say that $\mathcal{F}$ is *uniformly equicontinuous* [or *equiuniformly continuous*] on the subset A of Ω if

$$(\forall \epsilon > 0)(\exists \delta > 0)(\forall f \in \mathcal{F})(\forall a \in A) \quad y \in \Omega \ \& \ d_\Omega(a,y) < \delta \Longrightarrow d_\Psi(f(a), f(y)) < \epsilon .$$

Uniform continuity of f means that, for the given ϵ, we can choose a δ that "works" for f at all the points of A. In Example 1.7, f is not uniformly continuous on Ω, because the formula we found for the largest possible δ that works at x tends to 0 as $x \downarrow 0$. Whatever δ we tried would be too large if x were so small that $\epsilon x^2/(1+\epsilon x) < \delta$. Equiuniform continuity means that δ can be chosen to work for all the functions of $\mathcal{F}$ and all the points of A simultaneously. This certainly entails that $\mathcal{F}$ is equicontinuous at every point of A.

If $\mathcal{F}$ is a finite family, uniform equicontinuity is the same as uniform continuity of each function belonging to $\mathcal{F}$.

The family $\{f_\alpha : 0 < \alpha < \infty\}$ mentioned in Example 1.9 is non-equicontinuous at every point of $(0,1)$; but $\{f_\alpha : 0 < \alpha < 1\}$ is equicontinuous at every point of $(0,1)$, although none of the functions f_α is uniformly continuous on $(0,1)$.

On the other hand, if the formula (4) is used to define a family of functions $f_\alpha : (0,\infty) \longrightarrow \mathbb{R}$ for $0 < \alpha < 1$, this family is uniformly equicontinuous on $(1,\infty)$. But even the individual functions are not uniformly continuous on $(0,1)$.

The word "uniformly" also occurs in other contexts.

Definition 1.11. Suppose given a sequence (f_n) of functions $f_n : \Omega \longrightarrow \Psi$. We say that (f_n) *converges uniformly* to the function $f : \Omega \longrightarrow \Psi$, $f_n \to f$ uniformly, if

$$(\forall \epsilon > 0)(\exists N \in \mathbb{N})(\forall x \in \Omega) \quad n \geq N \Longrightarrow d_\Psi(f_n(x), f(x)) < \epsilon .$$

Then, for each fixed $x \in \Omega$, the sequence $(f_n(x))$ in Ψ tends to the limit $f(x)$; and there is a "rate of convergence" which applies independently of the choice of point x. [Stokes, when he introduced this idea in 1848, spoke of non-uniform convergence as "infinitely slow convergence", having in mind that there is no "rate of convergence" that governs the convergence at all points simultaneously. It is worth remarking that before the 20th century there was no significant division between pure and applied mathematics; Stokes is most famous for his work in fluid mechanics, although he made many other important contributions.]

Remark 1.12. In the definitions above we have consistently used "strong inequalities" $<$. The substance of each definition, for instance the *existence* of a positive δ for a given positive ϵ, is unchanged if some or all of the $<$ s are changed to $\leq$, though the δ in question for given ϵ may have to be altered; similarly, "$n \geq N$" could be changed to "$n > N$" without changing the sense of the definitions, because $n \geq N+1$ is equivalent to $n > N$. (However, we cannot change "$\epsilon > 0$" or "$\delta > 0$", since it is essential that they should be positive numbers.)

As far as we are concerned, the most important reason for considering uniform convergence is the following

Theorem 1.13. *Suppose that* (f_n) *is a sequence of functions* $\Omega \longrightarrow \Psi$ *that converges uniformly to* $f : \Omega \longrightarrow \Psi$*, and that each function* f_n *is continuous at* $x \in \Omega$*. Then* f *is also continuous at* $x \in \Omega$*.*

Proof. Let $\epsilon > 0$. There exists N such that, for every $y \in \Omega$, $d_\Psi(f_n(y), f(y)) < \frac{1}{3}\epsilon$ whenever $n \geq N$. But f_N is continuous at x; thus there is $\delta > 0$ such that

$$d_\Omega(x,y) < \delta \Longrightarrow d_\Psi(f_N(x), f_N(y)) < \tfrac{1}{3}\epsilon\,,$$

and it follows that, if $d_\Omega(x,y) < \delta$,

$$\begin{aligned} d_\Psi(f(x), f(y)) &\leq d_\Psi(f(x), f_N(x)) + d_\Psi(f_N(x), f_N(y)) + d_\Psi(f_N(y), f(y)) \\ &< \tfrac{1}{3}\epsilon + \tfrac{1}{3}\epsilon + \tfrac{1}{3}\epsilon = \epsilon\,, \end{aligned}$$

so that we have found a δ for the function f and the given ϵ at the point x. □

The reason for demanding *uniform* convergence is that we need $d_\Psi(f_N(y), f(y)) < \frac{1}{3}\epsilon$ for all points y within distance δ of x. Uniform convergence is a rather stronger condition than this, but it is certainly sufficient.

Corollary 1.14. *Suppose in* 1.15 *that each function* f_n *is uniformly continuous on* Ω*. Then so is* f*.* □

Indeed, δ may be chosen (for f_N) independently of x, and the proof is unchanged. We must assume uniform continuity of *each* f_n, because the proof exploits a particular f_N and we do not *a priori* know which; that depends on ϵ.

These results are examples of the "double limit problem", which arises in many contexts. Crudely put, one may be able to take two limits in succession, but if one tries to take the limits in the other order, it may not be possible or one may get a different answer. In 1.15, we want to show that, for any $x_n \to x$, $f(x_n) \to f(x)$; that is continuity at x. The information is that, for each fixed n, $f_m(x_n) \to f(x_n)$ as $m \to \infty$. Thus, by 1.6, we wish to prove that

$$\lim_{n\to\infty} \lim_{m\to\infty} f_m(x_n) = \lim_{m\to\infty} f_m(x) = \lim_{m\to\infty} \lim_{n\to\infty} f_m(x_n)$$

(since each f_m is continuous at x). The uniform convergence of the sequence of functions enables us to prove this, esentially by freeing one limit from dependence on the other. But we could also do it the other way round and get the less useful, but still interesting

Theorem 1.15. *Suppose that* (f_n) *is a sequence of functions* $\Omega \longrightarrow \Psi$ *that is equicontinuous at* $x \in \Omega$ [*that is, the family* $\{f_n : n \in \mathbb{N}\}$ *is equicontinuous at* x] *and such that, for each* $y \in \Omega$, $f_n(y) \to f(y)$ *as* $n \to \infty$*. Then* f *is also continuous at* $x \in \Omega$*.* □

§2. Complete metric spaces

Definition 2.1. Let (Ω, d) be a metric space, and $(x_n)_{n\in\mathbb{N}}$ a sequence therein. The sequence is *Cauchy* if

$$(\forall\epsilon > 0)(\exists N \in \mathbb{N}) \quad m, n \geq N \Longrightarrow d(x_m, x_n) < \epsilon\,.$$

Statements of this kind are often summarized as "$a_{mn} \to 0$ as $m, n \to \infty$", where in this case $a_{mn} := d(x_m, x_n)$.

For quite a while now the absurdity of using Cauchy's surname (or Banach's or Hilbert's) as an adjective has ceased to be felt, but previously "Cauchy sequences" were often called "fundamental", and this usage may still be found in old books.

Lemma 2.2. *A convergent sequence in a metric space Ω is automatically Cauchy.*

Proof. Suppose that $x_n \to x$, and take any $\epsilon > 0$. Then

$$(\exists N \in \mathbb{N})\; n \geq N \Longrightarrow d(x_n, x) < \tfrac{1}{2}\epsilon$$

(taking "$\frac{1}{2}\epsilon$" instead of "ϵ" in the definition of convergence). But now, if both $m \geq N$ and $n \geq N$,

$$d(x_m, x_n) \leq d(x_m, x) + d(x_n, x) < \tfrac{1}{2}\epsilon + \tfrac{1}{2}\epsilon = \epsilon\,. \qquad \square$$

Definition 2.3. The metric space (Ω, d) is (metrically) *complete* if every Cauchy sequence in Ω converges.

Thus, in a complete metric space, "Cauchy" and "convergent" are equivalent. The word "complete" is overused in mathematics, but we shall need only this sense.

It is one of the most important properties of $\mathbb{R}$ ("Cauchy's General Principle of Convergence", never actually *proved* by Cauchy) that, if you give it the standard ("Euclidean") metric $d(x, y) := |x - y|$, it forms a complete metric space: every Cauchy sequence of real numbers has a limit. This basic result entails — often not at all trivially — the completeness of many other metric spaces that are used in analysis.

The significance of completeness in the larger scheme of things is that it is what you often need to construct something by successive approximations. The approximations will form a Cauchy sequence, and the completeness is needed to ensure that there is a limit. This idea was systematized for many applications by Banach with the concept of a "contraction mapping".

Definition 2.4. Let Ω be a metric space, and let $f : \Omega \longrightarrow \Omega$ be a function. f is a *contraction mapping* (or just *contraction*) if there exists a number κ such that $0 \leq \kappa < 1$ [notice κ is *less than* 1] and, for all $x, y \in \Omega$, $d(f(x), f(y)) \leq \kappa d(x, y)$. κ may be called the *contraction constant*. Another way of putting it is that f is a "Lipschitz mapping" with Lipschitz constant κ strictly less than 1.

Theorem 2.5. (The contraction mapping principle.) *Let Ω be a complete metric space, and $f : \Omega \longrightarrow \Omega$ a contraction mapping of Ω. Then f has a unique fixed point in Ω; that is, there is one and only one point $f \in \Omega$ such that $f(\phi) = \phi$.*

Proof. Take any $\phi_0 \in \Omega$, and define a sequence $(\phi_n)_{n\in\mathbb{N}}$ in Ω by applying f repeatedly:

$$\phi_1 := f(\phi_0)\,, \quad \phi_2 := f(\phi_1)\,, \quad \ldots\,, \quad \phi_{n+1} := f(\phi_n)\,, \quad \ldots\,.$$

(It would be natural to write $\phi_n = f^n(\phi_0)$.) Let $q := d(\phi_0, \phi_1) = d(\phi_0, f(\phi_0))$, and let κ be the contraction constant. Then, by induction,

$$d(\phi_n, \phi_{n+1}) = d(f(\phi_{n-1}), f(\phi_n)) \le \kappa d(\phi_{n-1}, \phi_n) \le \cdots \le \kappa^n d(\phi_0, \phi_1) = \kappa^n q$$

for any $n \in \mathbb{N}$. Hence, if $n = m + k > m$,

$$\begin{aligned} d(\phi_m, \phi_n) &\le d(\phi_m, \phi_{m+1}) + d(\phi_{m+1}, \phi_{m+2}) + \cdots + d(\phi_{m+k-1}, \phi_n) \\ &\le (\kappa^m + \kappa^{m+1} + \cdots + k^{m+k-1})q = \kappa^m \frac{1-\kappa^k}{1-\kappa} q \le \frac{\kappa^m}{1-\kappa} q\,. \end{aligned} \tag{5}$$

As $\kappa^m \to 0$ as $m \to \infty$, it follows that $d(\phi_m, \phi_n) \to 0$ as $m, n \to \infty$, i.e. the sequence (ϕ_n) is Cauchy. But Ω is complete, and so there exists $\phi \in \Omega$ such that $\phi_n \to \phi$.

We now show that x is a fixed point of f. For any $n \in \mathbb{N}$,

$$\begin{aligned} 0 \le d(\phi, f(\phi)) &\le d(\phi, \phi_n) + d(\phi_n, \phi_{n+1}) + d(f(\phi), \phi_{n+1}) \\ &= d(\phi, \phi_n) + d(\phi_n, \phi_{n+1}) + d(f(\phi), f(\phi_n)) \\ &\le d(\phi, \phi_n) + d(\phi_n, \phi_{n+1}) + \kappa d(\phi, \phi_n)\,, \end{aligned}$$

where each term tends to 0 as $n \to \infty$. So, necessarily, $d(\phi, f(\phi)) = 0$, and $\phi = f(\phi)$.

Finally, suppose that ψ is another fixed point of f. Then

$$d(\phi, \psi) = d(f(\phi), f(\psi)) \le \kappa d(\phi, \psi)\,, \quad (1-\kappa) d(\phi, \psi) \le 0\,,$$

which, as $1 - \kappa > 0$ and $d(\phi, \psi) \ge 0$, is only possible if $d(\phi, \psi) = 0$ and $\phi = \psi$. □

The hypothesis that f is a contraction mapping ensures both the existence and the uniqueness of a fixed point, and the proof shows that, for any choice of ϕ_0, the result of applying f repeatedly, $\phi_0, f(\phi_0), f^2(\phi_0), f^3(\phi_0), \ldots$ is a sequence that converges to that one and only fixed point ϕ. For existence alone, it would suffice to know that f is continuous and that, for *some* $\phi_0 \in \Omega$, the sequence of iterates $(f^n(\phi_0))$ converges. ****

Many "existence proofs" can be formulated as "fixed point theorems". As a striking example, I shall give the standard theorem on the existence of solutions of ordinary differential equations. The result can be proved for equations in which the functions take values in higher-dimensional spaces, but to keep the discussion short I shall consider only the simplest reasonable hypotheses and result.

Definition 2.6. Let G be an open subset of $\mathbb{R}^2$. A function $g : G \longrightarrow \mathbb{R}$ satisfies a *Lipschitz condition in the second variable* on G, with *Lipschitz constant* $K \ge 0$, if, for any two points (x, y_1) and (x, y_2) in G, $|g(x, y_1) - g(x, y_2)| \le K|y_1 - y_2|$. [The Lipschitz constant here may be any nonnegative number.]

Theorem 2.7. *Let G be an open subset of $\mathbb{R}^2$, let $(x_0, y_0) \in G$, and let $g : G \longrightarrow \mathbb{R}$ be continuous and satisfy a Lipschitz condition in the second variable on G. There exists $\delta > 0$ such that the equation*

$$\frac{dy}{dx} = g(x,y)\,, \quad \textit{with initial condition } \; y(x_0) = y_0\,, \tag{6}$$

has a unique solution for $x_0 - \delta < x < x_0 + \delta$.

By a *solution* of the differential equation for $x_0 - \delta < x < x_0 + \delta$, we mean a function $y : (x_0 - \delta, x_0 + \delta) \longrightarrow \mathbb{R}$ which is differentiable on $(x_0 - \delta, x_0 + \delta)$ and which, for each $x \in (x_0 - \delta, x_0 + \delta)$, satisfies the conditions $(x, y(x)) \in G$ and $y'(x) = f(x, y(x))$. That this solution is *unique* means that there is only one function satisfying these conditions.

Proof. As g is continuous on G, there must be a neighbourhood G_0 of (x_0, y_0) and a constant $H \geq 0$ such that $G_0 \subseteq G$ and $|g(x,y)| \leq H$ for all $(x,y) \in G_0$. Choose $h > 0$ to be so small that, simultaneously, $\kappa := hK < 1$ and

$$(x_0 - h, x_0 + h) \times [y_0 - hH, y_0 + hH] \subseteq G_0\,. \tag{7}$$

Let Ω be the set of all the *continuous* functions

$$\phi : (x_0 - h, x_0 + h) \longrightarrow [y_0 - hH, y_0 + hH]\,.$$

If $\phi, \psi \in \Omega$, then $\phi(x) - \psi(x)$ is certainly bounded for $x_0 - h < x < x_0 + h$ (it takes values in $[-2hH, 2hH]$). Define

$$d(\phi, \psi) := \sup\{|\phi(x) - \psi(x)| : x_0 - h < x < x_0 + h\}\,, \tag{8}$$

and then d is a metric on Ω. Furthermore, (Ω, d) is complete. [I shall prove this as a separate lemma below.] ****

Suppose $\phi \in \Omega$. Define a function $\psi : (x_0 - h, x_0 + h) \longrightarrow [y_0 - hH, y_0 + hH]$ by

$$\psi(x) := y_0 + \int_{x_0}^{x} g(t, \phi(t))\, dt\,. \tag{9}$$

In the first place, the integral makes sense, since $\phi(t)$ is continuous in t and, therefore, $g(t, \phi(t))$ is also continuous in t. Secondly, our choice of h at (7), granted that ϕ is in Ω, ensures that the integrand has absolute value not exceeding H; so, for $x \in (x_0 - h, x_0 + h)$,

$$|\psi(x) - y_0| = \left| \int_{x_0}^{x} g(t, \phi(t))\, dt \right| \leq H|x - x_0| \leq hH\,,$$

whilst ψ is of course continuous, and even differentiable, in x, from (9). Thus $\psi \in \Omega$.

Hence, for each $\phi \in \Omega$, we have defined a corresponding $\psi \in \Omega$ by (9). We define $f : \Omega \longrightarrow \Omega$ by setting $f(\phi) := \psi$, as given by (9). f is thus a mapping from the "space of functions" Ω into itself.

But Ω was also a metric space with metric d (cf. (8)). With respect to this metric, I assert that f is a contraction mapping, with contraction constant κ. Given $x \in (x_0 - h, x_0 + h)$, and $\phi, \psi \in \Omega$,

$$\begin{aligned} |f(\phi)(x)) - f(\psi)(x)| &= \left| \left(y_0 + \int_{x_0}^{x} (g(t, \phi(t))\, dt \right) - \left(y_0 + \int_{x_0}^{x} g(t, \psi(t)))\, dt \right) \right| \\ &\leq \left| \int_{x_0}^{x} |g(t, \phi(t)) - g(t, \psi(t))|\, dt \right| \end{aligned}$$

$$\begin{aligned}&\le \left|\int_{x_0}^{x} K|\phi(t)) - \psi(t)|\,dt\right| \\ &\le Khd(\phi,\psi) = \kappa d(\phi,\psi)\,.\end{aligned}$$

(The first inequality is the "fundamental estimate" for integrals; the second uses the Lipschitz condition, and the large modulus signs are inserted in each because of the possibility that $x < x_0$.) Consequently, granted that (Ω, d) is complete, 2.5 shows that there exists a unique function $\chi \in \Omega$ such that $f(\chi) = \chi$. This means, by (9), that, for $x \in (x_0 - h, x_0 + h)$,

$$\chi(x) = y_0 + \int_{x_0}^{x} g(t, \chi(t))\,dt\,, \tag{10}$$

and the fundamental theorem of calculus tells us that in turn this is *equivalent* to

$$\frac{d\chi}{dx} = g(x, \chi(x)) \quad \text{for } |x - x_0| < h\,, \text{ and} \quad \chi(x_0) = y_0\,. \qquad \square$$

Lemma 2.8. (Ω, d) *is a complete metric space.*

Proof. Suppose, in fact, that (ϕ_n) is a Cauchy sequence in (Ω, d). Then, for any $\epsilon > 0$, there is some N (let us call it $N(\epsilon)$) such that, whenever $m, n \ge N(\epsilon)$,

$$\sup\{|\phi_m(x) - \phi_n(x)| : x_0 - h < x < x_0 + h\} < \epsilon$$

(recall (8)). Certainly, then, for any specific $x \in (x_0 - h, x_0 + h)$,

$$m, n \ge N(\epsilon) \Longrightarrow |\phi_m(x) - \phi_n(x)| < \epsilon\,, \tag{11}$$

and the sequence $(\phi_n(x))$ in $\mathbb{R}$ is a *numerical* Cauchy sequence. Therefore it has a limit in $\mathbb{R}$, which we may denote as $\phi(x)$ (in principle, of course, there will be a different limit for each choice of x). In this way we get a function $\phi : (x_0 - h, x_0 + h) \longrightarrow \mathbb{R}$.

Now, however, look at (11). If we "hold n constant, but let $m \to \infty$", we deduce

$$n \ge N(\epsilon) \Longrightarrow |\phi(x) - \phi_n(x)| \le \epsilon\,. \tag{12}$$

[*The proof of this: suppose that, for some (specific) $n \ge N(\epsilon)$, $|\phi(x) - \phi_n(x)| > \epsilon$. Then there exists M such that $m \ge M \Longrightarrow |\phi(x) - \phi_m(x)| < |\phi(x) - \phi_n(x)| - \epsilon$, this being a fixed positive number. But then, if $m \ge \max(M, N(\epsilon))$,

$$|\phi_m(x) - \phi_n(x)| \ge |\phi(x) - \phi_n(x)| - |\phi(x) - \phi_m(x)| > \epsilon\,,$$

which contradicts (11). Hence (12) must hold. Notice, however, that one can only obtain the weak inequality $\le$ in (12), despite the $<$ in (11). Simple examples show that this is inevitable.*] This is true for each $x \in (x_0 - h, x_0 + h)$, so that (compare 1.12) $\phi_n \to \phi$ uniformly on $(x_0 - h, x_0 + h)$. By 1.13, then, ϕ is continuous on $(x_0 - h, x_0 + h)$.

However, as, for each $x \in (x_0 - h, x_0 + h)$, $|\phi_n(x) - y_0| \le hH$, and $\phi_n(x) \to \phi(x)$, then $|\phi(x) - y_0| \le hH$ too (as in the proof of (12) between the asterisks). As ϕ is continuous and takes values in $[y_0 - hH, y_0 + hH]$, it belongs to Ω.

Finally, notice that, as (12) holds for any $x \in (x_0 - h, x_0 + h)$, we have

$$(\forall \epsilon > 0)(\exists N \in \mathbb{N}) \quad n \ge N \Longrightarrow \sup\{|\phi(x) - \phi_n(x)| : x_0 - h < x < x_0 + h\} \le \epsilon\,,$$

because N here may just be the $N(\epsilon)$ we had before. But this says that $d(\phi, \phi_n) \to 0$, or that $\phi_n \to \phi$ in the sense of the metric space (Ω, d). ☐

These proofs deserve several informal remarks.

The first is a general mathematical observation. In order to prove a result about a *differential* equation (6), we transformed it into the *integral* equation (10). This is a common trick in the theory of differential equations, and expresses the vague but important fact that the derivative of a differentiable function may be very irregular, for instance by oscillating very rapidly, whilst the function itself appears to behave quite regularly; or, to put it the other way round, that the indefinite integral of a function is "smoother" than the function itself. Integral equations of various kinds are, therefore, easier to handle than differential equations.

Secondly, the proof above is, in fact, not quite the original argument of Picard. He did not use the "contraction mapping principle" as such — it was invented thirty years or so later. But he defined f as above, and iterated $f^n(\phi_0)$ as in the proof of 2.5, starting with the function ϕ_0 that has constant value y_0. The Lipschitz condition (as long as it holds) can be used to show that the successive iterates converge, without restricting the size of h as we did (by requiring $hK < 1$) to ensure f is a contraction mapping. In this respect Picard's proof is somewhat better than ours, but the extra information it gives can easily be obtained from our weaker version; furthermore, our proof gives uniqueness at the same time, whilst Picard needed a separate (easy) argument.

Thirdly, the question naturally arises: how essential is the Lipschitz condition? It was clearly very important in our proof, although there are weaker but messier conditions that could be employed to much the same effect. The situation is, in fact, a little curious. There are simple examples in which the function g does not satisfy a Lipschitz condition and there is more than one solution of (6). But if g is merely *continuous* (and takes values in $\mathbb{R}^n$ for a finite dimension n) then (6) has *a* solution. This is a theorem of Peano. The proof is more difficult (the version usually given these days ultimately requires Brouwer's fixed-point theorem), but it also amounts to establishing that the mapping f has a fixed point.

Fourthly, the proof of Lemma 2.8 exemplifies my earlier remark that the completeness of $\mathbb{R}$ lies behind the completeness of many other useful metric spaces. My standard joke here (I have made it every year for 30 years) is that the mode of proof is rather like the procedure of an American political party during a presidential election. We are given a Cauchy sequence; we wish to show that it converges, and for this we must first of all guess what the limit is to be. We select our *candidate* for the limit by any reasonable procedure you can think of (in our case, we constructed ϕ as the "pointwise limit" of the ϕ_n, i.e. by taking the limit of the numerical sequence $(\phi_n(x))$ for each point x separately). They go for the person they think most electable, subject to compatibility with their guiding ideas. In both cases, the candidate proposed may not even belong to the right space (both Democrats and Republicans wanted Eisenhower to be their candidate, but his preference was unknown). Here, we had to show that our putative limit was indeed a continuous function with values in $[y_0 - hH, y_0 + hH]$, i.e. was really in Ω. And finally, you have to get the candidate elected; for us, this meant we had to prove that ϕ was the limit in the sense of the metric d, despite its construction as a far weaker sort of limit.

Some other applications of the contraction mapping principle are given in Kolmogorov and Fomin. Standard ones are the inverse mapping theorem and the Newton-Raphson method for finding numerical solutions of equations.

§3. Compactness

Definition 3.1. Let Ω be any set. A *topology* on, or in, Ω is a class $\mathcal{G}$ of subsets of Ω having the following properties.

(a) $\emptyset \in \mathcal{G}$ and $\Omega \in \mathcal{G}$.

(b) If $U, V \in \mathcal{G}$, then $U \cap V \in \mathcal{G}$.

(c) If $\mathcal{U}$ is any subclass of $\mathcal{G}$, i.e. $\mathcal{U} \subseteq \mathcal{G}$, then $\bigcup_{U \in \mathcal{U}} U \in \mathcal{G}$ too.

[Notice that the *members* of $\mathcal{G}$ are *subsets* of Ω, not points of Ω. To emphasize this, I have described $\mathcal{G}$ as a "class" of subsets rather than a "set" of subsets. No logical distinction between "classes" and "sets" is intended; it is just that $\mathcal{G}$ is, as it were, a "set" on a higher floor than Ω. Notice too that the subclass $\mathcal{U}$ in *(c)* may have very many members, even uncountably many.] The members of $\mathcal{G}$ are commonly described as the *open* sets of Ω with respect to the topology $\mathcal{G}$, so that $\mathcal{G}$ itself is often not named explicitly — it is simply the class of open sets, where $\emptyset$ and Ω must be open, the intersection of two open sets must be open, and any union of open sets must be open.

A subset of Ω is *closed* (with respect to $\mathcal{G}$) if its complement $\Omega \setminus A$ is open.

One may wish to consider several topologies on Ω; however, if one topology in particular is to be understood, one may speak simply of the *topological space* Ω without referring explicitly to the topology itself. Otherwise, I may write "the topological space $(\Omega, \mathcal{G})$".

Definition 3.2. Let A be any subset of the topological space Ω. The class $\mathcal{O}(A)$ of all open sets included in A (sets open in Ω that are subsets of A) is nonempty, since $\emptyset$ is one such. By 3.1*(c)*, $\bigcup_{U \in \mathcal{O}(A)} U$ is also open, and is also, of course, a subset of A. It is obviously the largest open subset of A (granted that the construction shows there is a largest open subset), and is called the *interior of* A (with respect to the given topology on Ω), $\mathrm{int}(A)$.

Similarly, the class $\mathcal{C}(A)$ of all closed sets including A is nonempty, as Ω is one, and $\bigcap_{C \in \mathcal{C}(A)} C$ is also closed (by applying 3.1*(c)* to complements). It is, therefore, the smallest closed set including A, which is called the *closure* of A in the given topology, $\mathrm{cl}(A)$. [I may write $\mathrm{cl}_\Omega(A)$ to emphasize that it is the closure in the topological space Ω, and so on.] Notice that A is closed if and only if $A = \mathrm{cl}(A)$ and that, for any A, $\mathrm{cl}(\mathrm{cl}(A)) = \mathrm{cl}(A)$.

These definitions say, almost explicitly, that $\mathrm{cl}(A) = \Omega \setminus (\mathrm{int}(\Omega \setminus A))$. It is quite possible that the interior of A may be $\emptyset$ and its closure may be Ω; for instance, in the topological space $\mathbb{R}$, the subset $\mathbb{Q}$ has null interior and its closure is $\mathbb{R}$.

Lemma 3.3. *Let Ω be a topological space, $A \subseteq \Omega$. Then $x \in \mathrm{cl}(A)$ if and only if every open set that contains x meets A. Thus any open set that meets $\mathrm{cl}(A)$ also meets A.* □

Definition 3.4. Let $f : \Omega \longrightarrow \Psi$ be a mapping between topological spaces. f is *continuous* if, for any open set U in the topology on Ψ, $f^{-1}(U) := \{x \in \Omega : f(x) \in U\}$ is open in the topology on Ω.

Remark 3.5. $f : \Omega_1 \longrightarrow \Omega_2$ is continuous if and only if, for any closed subset C of Ω_2, $f^{-1}(C)$ is closed in Ω_1. Indeed, for any $S \subseteq \Omega_2$, $f^{-1}(\Omega_2 \setminus S) = \Omega_1 \setminus f^{-1}(S)$.

Definition 3.6. Let (Ω, d) be a metric space. A set $A \subseteq \Omega$ is *metrically open* if, for any $a \in A$, there is some $r > 0$ such that the *open metric ball* of radius r about a,

$$B(a;r) \coloneqq \{x \in \Omega : d(a,x) < r\} \subseteq A\,.$$

The class of all the metrically open subsets of Ω is a topology on Ω [this was proved in 312, and is not difficult], called the *metric topology*. Thus the metrically open sets are the open sets of the metric topology.

In 312, I showed that open metric balls are themselves open in the metric topology, and that a mapping between metric spaces is continuous in the sense of 3.4 with respect to the metric topologies if and only if it is continuous at every point $x \in \Omega$ in the sense of 1.5. From a logical point of view, though not practically, 3.4 is obviously a simpler definition.

If Ω is a metric space and $A \subseteq \Omega$, $\mathrm{int}(A)$ is the set of points $x \in \Omega$ such that $B(x;r) \subseteq A$ for some sufficiently small positive number r, and $\mathrm{cl}(A)$ is the set of points of Ω that are limits (in Ω) of sequences in A. An alternative characterization of $\mathrm{cl}(A)$ is that

$$x \in \mathrm{cl}(A) \text{ if and only if, for all } \epsilon > 0\,,\ B(x;\epsilon) \cap A \neq \emptyset\,.$$

Lemma 3.7. *A set X is closed in the metric topology on Ω if and only if the limit of a convergent sequence (x_n) in Ω whose terms all belong to X also belongs to X.* □

Lemma 3.8. *A subset of a complete metric space Ω is a complete metric space in the subspace metric if and if it is closed in Ω.* □

In this course I shall mostly be concerned with metric spaces, but a little information on topological spaces will help occasionally. The properties we deal with are often topological ones — that is, they can be expressed entirely in terms of the topology.

Definition 3.9. Let A be a subset of the topological space $(\Omega, \mathcal{G})$. Then A is said to be *compact* (or, in older books, to have the *Heine-Borel property*) with respect to $\mathcal{G}$ if any class $\mathcal{U}$ of open sets such that $A \subseteq \bigcup_{U \in \mathcal{U}} U$ [this is customarily expressed by saying that $\mathcal{U}$ is an *open covering of* A] includes a subclass $\mathcal{V}$ such that $A \subseteq \bigcup_{U \in \mathcal{V}} U$ [$\mathcal{V}$ is a *subcovering* of the covering $\mathcal{U}$ of A] and $\mathcal{V}$ has only finitely many members. ["Every open covering of A has a finite subcovering".]

Definition 3.10. Let A be a subset of the metric space (Ω, d). Then A is *sequentially compact* if every sequence (a_n) in A [each term a_n is a point of A] has a subsequence that converges to a point of a.

An important theorem, proved at some length in 312, is the following.

Theorem 3.11. *A subset A of the metric space Ω is sequentially compact if and only if it is compact with respect to the metric topology.* □

Thus sequential compactness, defined originally in terms of the metric, will remain true if the metric is changed to another that defines the same metric topology, so that for metric spaces there is no reason to include the adjective "sequential". This is a typical instance of the way compactness is (despite first impressions) a "natural" property. As I said in 312, it is a sort of "topological analogue of finiteness". Finite sets are obviously compact; but the remarkable thing is that many more complicated sets are too.

Lemma 3.12. *A compact metric space is complete.* □

Theorem 3.13. *A subset A of $\mathbb{R}^n$ is compact if and only if it is closed and bounded.* $\square$

"Bounded" here means that the set $\{d(a,0) : a \in A\}$ is bounded above in $\mathbb{R}$ (where d is of course the Euclidean metric).

Theorem 3.14. *Let $f : \Omega \longrightarrow \Psi$ be a continuous mapping between topological spaces. If K is a compact subset of Ω, $f(\Omega)$ is a compact subset of Ψ.* $\square$

Corollary 3.15. *If K is a compact subset of a topological space Ω, and $f : \Omega \longrightarrow \mathbb{K}$ is a continuous function, where $\mathbb{K}$ denotes either $\mathbb{R}$ or $\mathbb{C}$ with the standard metric, then $f(K)$ is a bounded closed subset of $\mathbb{K}$, and so*

(a) *if $\mathbb{K} = \mathbb{R}$, there is a point $k \in K$ such that $f(k) = \sup\{f(x) : x \in K\} < \infty$,*
(b) *in either case, there is $k \in K$ such that $|f(k)| = \sup\{|f(x)| : x \in K\} < \infty$.* $\square$

["sup A" denotes the *supremum* or *least upper bound* of the set $A \subseteq \mathbb{R}$].

Lemma 3.16. (a) *Let A be a compact subset and C be a closed subset of the topological space Ω. Then $A \cap C$ is compact.*

(b) *Let A be a compact subset of Ω, and suppose $\mathcal{C}$ is a class of closed sets such that $A \cap C \neq \emptyset$ for each $C \in \mathcal{C}$, and such that, whenever $C_1, C_2 \in \mathcal{C}$, either $C_1 \subseteq C_2$ or $C_2 \subseteq C_1$. Then $A \cap (\bigcap_{C\in\mathcal{C}} C) \neq \emptyset$.*

Proof. *(a)* Let $\mathcal{U}$ be an open cover of $A \cap C$. Then $\mathcal{U} \cup \{\Omega \setminus C\}$ is an open cover of A, so there is a finite subcover $\{U_1, U_2, \ldots, U_k, \Omega \setminus C\}$ of A. But, as $(\Omega \setminus C) \cap (A \cap C) = \emptyset$, $A \cap C$ must be covered by $\{U_1, U_2, \ldots, U_k\}$, which is an open subcover of $\mathcal{U}$.

(b) Suppose $A \cap (\bigcap_{C\in\mathcal{C}} C) = \emptyset$. Then $\{\Omega \setminus C : C \in \mathcal{C}\}$ is an open cover of A, so there is a finite subcover $\{\Omega \setminus C_1, \Omega \setminus C_2, \ldots, \Omega \setminus C_n\}$. One of $C_1, C_2, \ldots, C_n$ is smallest (compare C_1 and C_2; then compare the smaller of them with C_3; and so on); suppose, for example, that it is C_n. Then $\Omega \setminus C_n$ includes all the other sets in the cover, $A \subseteq \Omega \setminus C_n$, which implies $A \cap C_n = \emptyset$; and this is a contradiction. $\square$

Definition 3.17. Let Ω be a topological space. It is *locally compact* if, for any $x \in \Omega$ and any open set U in Ω such that $x \in \Omega$, there are an open set V and a compact set K in Ω such that $x \in V \subseteq K \subseteq U$.

Ω is *regular* if, for any $x \in \Omega$ and any open set U in Ω such that $x \in \Omega$, there are an open set V and a closed set K in Ω such that $x \in V \subseteq K \subseteq U$.

[Thus, in a regular space, each point x has a "base of closed neighbourhoods", and in a locally compact space, each point has a "base of compact neighbourhoods". These are the standard definitions, but Kelley in his well-known book requires for local compactness only *one* compact neighbourhood of each point. He then uses regular spaces with this property, which, because of the lemma that follows, are also locally compact in our stronger sense.]

Lemma 3.18. *Let Ω be a regular space, and $x \in \Omega$. Suppose there are an open set W and a compact set K such that $x \in W \subseteq K$. For any open set $U \ni x$, there are a closed compact set C and an open set V such that $x \in V \subseteq C \subseteq U$.*

Proof. $U \cap W$ is an open set containing x, so, by regularity, there are a closed set C and an open set V such that $x \in V \subseteq C \subseteq U \cap W$. But 3.16*(a)* then implies that $C = C \cap K$ is compact as well as closed. □

Definition 3.19. Let $f : \Omega \longrightarrow \Psi$ be a mapping between topological spaces. It is *open* if, for any open set U in Ω, its image $f(U)$ under f is open in Ψ.

Notice, firstly, that "openness" of a mapping is different from "openness" of a set. As so often, the context decides what meaning is intended. Secondly, if f is both one-one and onto, the inverse mapping f^{-1} that exists in that case is continuous if and only if f is open; and f is continuous if and only if f^{-1} is open. But, in general, continuity and openness of a mapping are unrelated.

Definition 3.20. Let Ω and Ψ be topological spaces. If V is an open set in Ω and W is an open set in Ψ, then $V \times W$ may be described as a *product open set* in $\Omega \times \Psi$. A subset of $\Omega \times \Psi$ is *open in the product topology* if it is a union of product open sets. Thus, a set $Z \subseteq \Omega \times \Psi$ is open in the product topology if and only if, for any $(x, y) \in Z$, there are open sets V in Ω and W in Ψ such that $x \in V$, $y \in W$, and $V \times W \subseteq Z$. The *product topology* itself is the class of subsets of $\Omega \times \Psi$ that are "open in the product topology".

Now suppose the topology of Ω is defined by a metric d_Ω and the topology of Ψ is defined by a metric d_Ψ. The product topology on $\Omega \times \Psi$ is defined by the metric d, where

$$d((x_1, y_1), (x_2, y_2)) = d_\Omega(x_1, x_2) + d_\Psi(y_1, y_2) \tag{13}$$

for $x_1, x_2 \in \Omega$ and $y_1, y_2 \in \Psi$.

The product topology may also be defined by many other metrics on $\Omega \times \Psi$, and the formula (13) has no claim to define *the* "product metric". ****

§4. Normed spaces

Definition 4.1. Let E be a vector space over $\mathbb{K}$, where $\mathbb{K}$ denotes either $\mathbb{R}$ or $\mathbb{C}$. A *norm* in E is a function $\nu : E \longrightarrow \mathbb{R}$ such that, for any $x, y \in E$ and $\lambda \in \mathbb{K}$,

$$\begin{array}{ll} \textit{(a)} & \nu(x + y) \le \nu(x) + \nu(y)\,, \\ \textit{(b)} & \nu(\lambda x) = |\lambda|\, \nu(x)\,, \\ \textit{(c)} & \nu(x) = 0 \Longrightarrow x = 0\,. \end{array}$$

Notice that *(a)* and *(b)* force ν to take nonnegative values. *(a)* is expressed by saying that the norm is *subadditive*, *(b)* by saying that it is *homogeneous*, and *(c)* by saying that it is *definite*. If *(a)* and *(b)* hold, we say ν is a *seminorm* (or pseudonorm). The whole package consisting of E with its vector space structure and the (semi)norm ν is called a (semi)*normed space*. It is rather usual to denote a norm by $\|\|$; that is, we write $\|x\|$ instead of $\nu(x)$.

A (semi)normed space $(E, \|\|)$ is a (pseudo)metric space if we define a metric d by

$$d(x, y) := \|x - y\|\,. \tag{14}$$

Only *(a)* and *(c)* are needed for this. Thus we can apply all the ideas and terminology of metric spaces (convergence of sequences, continuity of functions, compactness) to normed spaces; but, since the metric we have in mind is so closely related to the norm, we do not usually bother to introduce a separate notation for it.

Definition 4.2. The normed space E is a *Banach space* if it is complete as a metric space.

Definition 4.3. Let Ω be a compact topological space (i.e. Ω is a compact subset of itself). Then $C_{\mathbb{K}}(\Omega)$, or $C(\Omega;\mathbb{K})$, or $C(\Omega)$ if no ambiguity is possible, denotes the set of all continuous mappings $\Omega \longrightarrow \mathbb{K}$.

It is customary to endow $C(\Omega;\mathbb{K})$ with the structure of a vector space over the field $\mathbb{K}$, by taking the linear operations "pointwise":

$$(\forall\lambda,\mu \in \mathbb{K})(\forall f,g \in C(\Omega;\mathbb{K}))(\forall x \in \Omega) \quad (\lambda f + \mu g)(x) := \lambda(f(x)) + \mu(g(x))\,.$$

Furthermore, the real or complex vector space $C(\Omega;\mathbb{K})$ may be given a norm, the so-called *supremum norm* or *uniform norm*. If $f \in C(\Omega;\mathbb{K})$, then, by 3.15, $\{|f(x)| : x \in \Omega\}$ is a bounded subset of $\mathbb{R}$. Set

$$\|f\| := \sup\{|f(x)| : x \in \Omega\}\,. \tag{15}$$

This is very easily seen to be a norm on $C(\Omega;\mathbb{K})$. *****

$C(\Omega;\mathbb{K})$ also has the structure of an (associative) *algebra* over $\mathbb{K}$. That is: one defines the (pointwise) *product* fg of $f \in C(\Omega)$ and $g \in C(\Omega;\mathbb{K})$ by

$$(\forall x \in \Omega) \quad (fg)(x) := f(x)g(x)\,.$$

(The right-hand side is the product in $\mathbb{K}$ of two numbers, $f(x)$ and $g(x)$, in $\mathbb{K}$.) It is trivial that this multiplication is bilinear and associative. Furthermore,

$$(\forall f,g \in C(\Omega;\mathbb{K})) \quad \|fg\| \le \|f\|\|g\|\,. \text{****}$$

This last property is often expressed by the statement that *the norm is submultiplicative*.

Lemma 4.4. $C(\Omega;\mathbb{K})$ *is a Banach space with respect to the norm* $\|\|$.

Proof. It is only the completeness that needs to be checked, and the argument is really the same as in 2.8, but, in the absence of a metric, has to look rather different.

Let (f_n) be a Cauchy sequence in $C(\Omega;\mathbb{K})$. Thus

$$(\forall\epsilon > 0)(\exists N(\epsilon)) \quad m,n \ge N(\epsilon) \Longrightarrow \|f_m - f_n\| < \epsilon\,,$$

which, in view of the definition of the norm as the least upper bound, implies that

$$(\forall\epsilon > 0)(\exists N(\epsilon))(\forall x \in \Omega) \quad m,n \ge N(\epsilon) \Longrightarrow |f_m(x) - f_n(x)| < \epsilon\,.$$

But this implies that, for any specific $x \in \Omega$,

$$m,n \ge N(\epsilon) \Longrightarrow |f_m(x) - f_n(x)| < \epsilon\,,$$

so that the numerical sequence $(f_n(x))$ is Cauchy in $\mathbb{K}$. Therefore it has a limit in $\mathbb{K}$. In principle this limit will vary when x changes; denote it by $f(x)$. f is thus a *function* $\Omega \longrightarrow \mathbb{K}$, the 'pointwise limit' of (f_n). However, letting $m \to \infty$, one finds that $n \ge N(\epsilon) \Longrightarrow |f(x) - f_n(x)| \le \epsilon$ for every $x \in \Omega$, and so (recall 1.11 and 1.12) $f_n \to f$ uniformly on Ω. It is necessary to show that f is continuous. If Ω is a metric space, we can appeal to 1.13; but the result that a uniform limit of continuous functions is continuous is true when Ω is merely a topological space. Here is a proof.

Let U be open in $\mathbb{K}$. We wish to show that $f^{-1}(U)$ is open. Let $u \in f^{-1}(U)$. Thus $f(u) \in U$, and there is some $\epsilon > 0$ such that $B_{\mathbb{K}}(f(u);\epsilon) \subseteq U$ (recall 3.6). Take $N := N(\frac{1}{3}\epsilon)$. Now $V(u) := f_N^{-1}(B_{\mathbb{K}}(f_N(u);\frac{1}{3}\epsilon))$ is open, because f_N is continuous and $B_{\mathbb{K}}(f(u);\epsilon)$ is open in $\mathbb{K}$, and it obviously contains u. Let $y \in V(u)$. Then

$$\begin{aligned}|f(y) - f(u)| &\leq |f(y) - f_N(y)| + |f_N(y) - f_N(u)| + |f_N(u) - f(u)| \\ &\leq \tfrac{1}{3}\epsilon + |f_N(y) - f_N(u)| + \tfrac{1}{3}\epsilon < \epsilon,\end{aligned}$$

the outer terms being dominated by $\sup\{|f(x) - f_N(x)| : x \in \Omega\}$ and the middle one restricted by the definition of $V(u)$. Hence, $f(y) \in B_{\mathbb{K}}(f(u);\epsilon)$, and $V(u) \subseteq f^{-1}(U)$. So, any point u of $f^{-1}(U)$ lies in an open set $V(u)$ included in $f^{-1}(U)$, and $f^{-1}(U)$ is open (as the union of all these open sets $V(u)$). This shows that f is continuous, $f \in C(\Omega;\mathbb{K})$.

Finally, $n \geq N(\epsilon) \Longrightarrow \|f - f_n\| = \sup\{|f(x) - f_n(x)| : x \in \Omega\} \leq \epsilon$. That is, $f_n \to f$ in the sense of the norm in $C(\Omega;\mathbb{K})$. □

I refer you to my American joke. f is the obvious candidate for a limit of (f_n), but it is not *a priori* clear that it is a point of $C(\Omega;\mathbb{K})$, that is, a continuous function $\Omega \longrightarrow \mathbb{K}$, or that it is the limit in the sense of the norm, since it was defined in a different way. Notice, too, the finicky logical point that we can't legitimately write $\|f - f_n\|$ until we have shown that $f \in C(\Omega;\mathbb{K})$; for we defined the norm only on $C(\Omega;\mathbb{K})$.

We shall have many other examples of Banach spaces later. But I hope this example convinces you that the concept is worth having. In fact $C(\Omega;\mathbb{K})$ is a Banach *algebra* over $\mathbb{K}$.

Definition 4.5. Suppose that A is a normed space over $\mathbb{K}$, with norm $\|\|$, and an (associative) algebra (with multiplication denoted by juxtaposition) over $\mathbb{K}$. It is described as a *normed algebra* if the multiplication $(a,b) \mapsto ab : A \times A \longrightarrow A$ is continuous when $A \times A$ is given the product topology, and as a *Banach algebra* if $(A, \|\|)$ is in addition a Banach space.

From 3.20, with (13) and (14), we see that the product topology on $A \times A$ may be defined by a metric

$$d((a,b),(x,y)) := \|a - x\| + \|b - y\|.$$

Lemma 4.6. *A is a normed algebra if and only if there is some $K \geq 0$ such that*

$$(\forall x, y \in A) \quad \|xy\| \leq K\|x\|\|y\|. \tag{16}$$

Proof. Suppose the condition is satisfied, and that $(x_n, y_n) \to (x, y)$ in $A \times A$. That means $\|x_n - x\| + \|y_n - y\| \to 0$, or, equivalently, that $\|x_n - x\| \to 0$ and $\|y_n - y\| \to 0$. We may omit the initial terms of the sequence and assume that $\|y_n - y\| < 1$ for all n. Then

$$\begin{aligned}\|x_n y_n - xy\| = \|(x_n - x)y_n + x(y_n - y)\| &\leq \|(x_n - x)y_n\| + \|x(y_n - y)\| \\ &\leq K\|x_n - x\|\|y_n\| + K\|x\|\|y_n - y\| \\ &\leq K\|x_n - x\|(1 + \|y\|) + K\|x\|\|y_n - y\| \to 0.\end{aligned}$$

Conversely, suppose the condition is *not* satisfied, so that there is no K satisfying (16). For each $n \in \mathbb{N}$, there must be some x_n, y_n such that $\|x_n y_n\| > n^2\|x_n\|\|y_n\|$. This can only be so if $x_n \neq 0 \neq y_n$, clearly. But now let $x'_n := \dfrac{x_n}{n\|x_n\|}$, $y'_n := \dfrac{y_n}{n\|y_n\|}$: as $\|x'_n\| = \frac{1}{n} = \|y'_n\|$, we have $(x'_n, y'_n) \to (0,0)$, but, for all n,

$$\|x'_n y'_n\| = \frac{\|x_n y_n\|}{n^2 \|x_n\| \|y_n\|} > 1,$$

so that $x'_n y'_n \not\to 0$, and the multiplication is not continuous. □

Of course, if (16) is true, we may suppose that $K > 0$. Define a new norm $|\!|\!| \; |\!|\!|$ on A by

$$(\forall x \in A) \quad |\!|\!| x |\!|\!| := K\|x\|.$$

Then $|\!|\!| \; |\!|\!|$ is a norm on A that defines the same topology as $\|\|$ (the same open sets, the same convergent sequences, and even the same Cauchy sequences), but (16) implies

$$(\forall x, y \in A) \quad |\!|\!| xy |\!|\!| \leq |\!|\!| x |\!|\!| \, |\!|\!| y |\!|\!|. \tag{17}$$

The relation (17) is expressed by saying that the norm $|\!|\!| \; |\!|\!|$ is *submultiplicative* (with respect to the algebra multiplication). You will find that people often define normed algebras by insisting that the norm be submultiplicative, and, indeed, most of the normed algebras that arise in practice have “naturally defined” norms that are submultiplicative to begin with — but that is not an essential part of the theory. In particular, the norm we defined in $C(\Omega; \mathbb{K})$ is trivially submultiplicative.

§5 Baire category

Metric completeness is what allows constructions by successive approximations. In the 1920s a number of startling results, soon to be discussed, were proved by this means, and in the early 1930s Saks pointed out that the several complicated constructions that had been used could be avoided, or at least consolidated into one argument, by appealing to some older ideas of Baire. Let me first give a few familiar definitions.

Definition 5.1. Let A, B be subsets of the topological space Ω. A is *dense in* B if $B \subseteq \mathrm{cl}(A)$. Equivalently, any closed set that includes A also includes B; or, any open set that meets B must also meet A. If Ω is a metric space, yet another way of expressing the same situation is that every point of B is the limit in Ω of a sequence of points of A.

A subset D of Ω is *nowhere dense* (in Ω) if $\emptyset$ is the only open set of Ω in which D is dense; that is, the closure of D has empty interior. Equivalently, the interior of the complement of D (the same as the complement of the closure of D) is dense in Ω. A set is nowhere dense if it is included in the complement of an open set that is dense in Ω.

Notice that the closure of a nowhere dense set is nowhere dense by definition, and that any subset of a nowhere dense set is nowhere dense.

A subset E of Ω is *of the first category of Baire* in Ω [or, sometimes, is *meagre*] if it may be expressed as a countable union of nowhere dense sets. Otherwise, it is *of the second category of Baire*. A subset F of Ω is *residual* if $\Omega \setminus F$ is of the first Baire category.

Notice that, from the definition, a countable union of sets of first category is itself of first category. At first sight this seems to suggest that the class of such sets must contain some rather “large” sets; it is not clear that the second category contains any non-empty sets at all.

Lemma 5.2. *Let* (Ω, d) *be a complete metric space, and suppose that* $(C(x_n;\epsilon_n))$ *is a decreasing sequence of closed metric balls, where* (ϵ_n) *is a sequence of positive numbers tending to* 0. *Then* $\bigcap_{m=1}^{\infty} C(x_m;\epsilon_m) \neq \emptyset$. *Indeed, it is a singleton.*

Proof. If $n \geq m$, then $x_n \in C(x_n;\epsilon_n) \subseteq C(x_m;\epsilon_m)$, so $d(x_n, x_m) \leq \epsilon_m$. Hence (x_n) is a Cauchy sequence. Since Ω is complete, $x_n \to y \in \Omega$. But, as $x_n \in C(x_m;\epsilon_m)$ for $n \geq m$, and $C(x_m;\epsilon_m)$ is closed, therefore $y \in C(x_m;\epsilon_m)$, and this is true for any choice of m: $y \in \bigcap_{m=1}^{\infty} C(x_m;\epsilon_m)$.

Suppose $y \neq z \in \bigcap_{m=1}^{\infty} C(x_m;\epsilon_m)$. There exists m such that $\epsilon_m < \frac{1}{2}d(y,z) > 0$, but this contradicts $d(y,x_m) \leq \epsilon_m \geq d(z,x_m)$. So $\bigcap_{m=1}^{\infty} C(x_m;\epsilon_m)$ must be a singleton. ☐

Proposition 5.3. *Let* U *be a non-empty open set in a complete metric space* (Ω, d). *For* $n = 1, 2, 3, \ldots$, *let* U_n *be open in* Ω *and dense in* U. *Then* $\bigcap_{n=1}^{\infty} U_n$ *is also dense in* U.

Proof. Let V be any open set that meets U. We wish to show (see 5.1) that V must also meet $\bigcap_{n=1}^{\infty} U_n$. Take any $x_0 \in U \cap V$, which is non-empty and open; there is, therefore, some $\epsilon_0 > 0$ such that $B(x_0;\epsilon) \subseteq U \cap V$. By hypothesis, U_1 is dense in U, and so $B(x_0;\frac{1}{2}\epsilon_0) \cap U_1 \neq \emptyset$; thus, there exists $x_1 \in B(x_0;\frac{1}{2}\epsilon_0) \cap U_1$, which is open. In turn, there is $\epsilon_1 > 0$ such that $\epsilon_1 \leq \frac{1}{2}\epsilon_0$ and $B(x_1;\epsilon_1) \subseteq B(x_0;\frac{1}{2}\epsilon_0) \cap U_1$.

Suppose that $x_0, x_1, \ldots, x_n, \epsilon_0, \epsilon_1, \ldots, \epsilon_n$ (where $n \geq 1$) have been chosen, and

$$\epsilon_k \leq \tfrac{1}{2}\epsilon_{k-1}, \quad B(x_k;\epsilon_k) \subseteq B(x_{k-1};\tfrac{1}{2}\epsilon_{k-1}) \cap U_k$$

for $1 \leq k \leq n$. Then (as $x_n \in B(x_0;\frac{1}{2}\epsilon_0) \subseteq U$, and U_{n+1} is dense in U) there is some $x_{n+1} \in B(x_n;\frac{1}{2}\epsilon_n) \cap U_{n+1}$, and in turn there is $\epsilon_{n+1} > 0$ such that $\epsilon_{n+1} \leq \frac{1}{2}\epsilon_n$ and $B(x_{n+1};\epsilon_{n+1}) \subseteq B(x_n;\frac{1}{2}\epsilon_n) \cap U_{n+1}$. The construction may be carried out inductively.

Now, for each n, $C(x_{n+1};\frac{1}{2}\epsilon_{n+1}) \subseteq B(x_{n+1};\epsilon_{n+1}) \subseteq B(x_n;\frac{1}{2}\epsilon_n) \subseteq C(x_n;\frac{1}{2}\epsilon_n) \subseteq U_n$, and of course $\epsilon_n \leq \frac{1}{2}\epsilon_{n-1} \leq \cdots \leq 2^{-n}\epsilon_0$, so that $\epsilon_n \to 0$. By 5.2, $\bigcap_{n=0}^{\infty} C(x_n;\frac{1}{2}\epsilon_n)$ contains a (unique) point y. However, $C(x_n;\frac{1}{2}\epsilon_n) \subseteq U_n$ for each n, so that $y \in \bigcap_{n=1}^{\infty} U_n$, and also $y \in B(x_0;\frac{1}{2}\epsilon_0) \subseteq V$, so that $V \cap (\bigcap_{n=1}^{\infty} U_n) \neq \emptyset$. ☐

Theorem 5.4. (The Baire category theorem.) *A non-empty open set in a complete metric space* Ω *is of the second category of Baire in* Ω.

Proof. Suppose, instead, that the non-empty open set U is a countable union $\bigcup_{n=1}^{\infty} D_n$ of nowhere dense sets. Thus $U \subseteq \bigcup_{n=1}^{\infty} \mathrm{cl}_\Omega(D_n)$, and $U \cap (\bigcap_{n=1}^{\infty}(\Omega \setminus \mathrm{cl}_\Omega(D_n))) = \emptyset$. But $\Omega \setminus \mathrm{cl}_\Omega(D_n)$ is open and dense in Ω for each n (this being another formulation of "nowhere dense"; see 5.1); hence, 5.3 gives a contradiction. ☐

This is the form in which the Baire category theorem is customarily stated, but 5.3 easily yields a little more. If we regard the open set U as a metric space in its own right, a subset of U is of the first category in U if and only if it is of the first category in Ω. **** In this sense, U is of the second category in U, as well as in Ω. That is trivial; but, more substantially, any set of the first category in U has a relative complement in U that is dense in U. The word "residual" is sometimes used to mean a set whose complement is of first category; thus in a space which is a non-empty open subset of a complete metric space, any residual set is dense. [Such a space will often not be a complete metric space in its own right.]

There is a second version of 5.4. Instead of 5.3, one may use 3.16*(b)* (the fact that the intersection in 5.3 is a *singleton* is not needed to deduce 5.4). The proof is in essence much the same, and the result is the following.

Theorem 5.5. *Let Ω be a regular locally compact topological space. Any non-empty open subset U of Ω is of the second category of Baire in Ω.* ***** □

In this case too, U will be of the second category of Baire in U itself. If I had discussed the subspace topology, the statement would be trivial, because U is also regular locally compact. The similar result for an open subset of a complete metric space was not so immediate, since the open subset is not itself complete (it is "locally complete", however, which is all that is really needed for 5.4).

5.3 constructs a point of $V \cap (\bigcap_{n=1}^{\infty} U_n)$ by successive approximation. (This cannot be said of 5.5, but it is a far less interesting result.) A rather non-trivial example of the way in which one may use the idea is the following.

Theorem 5.6. *The set of continuous functions $[0,1] \longrightarrow \mathbb{R}$ that are non-differentiable at every point of $[0,1]$ is residual in $C([0,1];\mathbb{R})$.*

This is one sense in which the "everywhere non-differentiable" functions are very much commoner than the others. I shall not give all the details of the proof.

Proof. Given $n, m \in \mathbb{N}$, let G_{nm} be the set of functions $f \in C([0,1];\mathbb{R})$ such that

$$2^n |f(2^{-n}k) - f(2^{-n}(k+1))| > 2^m \tag{18}$$

for $0 \leq k \leq 2^n - 1$. It is an open subset of $C([0,1];\mathbb{R})$. (Check this.)

Let $G_m := \bigcup_{n=1}^{\infty} G_{nm}$. This is also open, as a union of open sets. But it is also *dense* in $C([0,1];\mathbb{R})$. I shall only sketch the proof.

Take any $f \in C([0,1];\mathbb{R})$ and $\epsilon > 0$. Suppose $2^{-q} \leq \epsilon$. f is uniformly continuous; so there exists $N \in \mathbb{N}$ such that

$$|x - y| \leq 2^{-N} \Longrightarrow |f(x) - f(y)| < 2^{-q-1}.$$

Construct a function f_1 to agree with f at the points $2^{-N}k$, $0 \leq k \leq 2^N$, and to have a straight-line graph between adjacent points. Then $\|f_1 - f\| < 2^{-q-1}$ too, and the slopes of the straight-line segments will not exceed 2^{N-q-1} in absolute value.

Add to f_1 another piecewise-linear function g which oscillates, its vertices being at the points $(2^{-q-m-2-N}k, (-1)^k 2^{-q-1})$. The slopes of the segments here are $\pm 2^{N+m+1}$; thus the slopes of the segments of $f_2 := f_1 + g$ are at least 2^{m+1} in absolute value, and $f_2 \in G_{sm}$, where $s := q + m + 2 + N$. But $\|g\| = 2^{-q-1}$, and so $\|f_2 - f\| < 2^{-q} \leq \epsilon$. This proves that G_m is dense in $C([0,1];\mathbb{R})$, as asserted.

By the Baire category theorem, $\bigcap_{m=1}^{\infty} G_m$ is dense (in fact residual) in $C([0,1];\mathbb{R})$. However, suppose $f \in C([0,1];\mathbb{R})$ is differentiable at x. Then, as $n \to \infty$, the quotients (18) for which $2^{-n}k \leq x \leq 2^{-n}(k+1)$ tend to $f'(x)$, and so are bounded. This evidently implies that $f \notin \bigcap_{m=1}^{\infty} G_m$. So the set of functions that are everywhere non-differentiable includes the residual set $\bigcap_{m=1}^{\infty} G_m$, and is itself residual. □

Following the proof of 5.4 step by step, one may construct specific examples of everywhere non-differentiable continuous functions (such as the familiar example of Weierstrass), but the argument from Baire category gives a defensible form to the vague general statement that, amongst continuous functions, the nowhere differentiable continuous ones are in a sense "usual", and differentiability even at a single point is "unusual".

There is a more precise sense in which this is true. The functions that are somewhere differentiable form a set of Wiener measure zero in $C([0,1];\mathbb{R})$; in effect, there is zero probability, in a fairly natural sense, that a continuous function on $[0,1]$ will have even one point of differentiability. [In general, there is a close analogy between "measure zero" and "first category", but they are different concepts. There are sets of first category in $[0,1]$ that have Lebesgue measure 1; their complements will be of second category and of measure 0.

§6. The Stone-Weierstraß theorem.

The result of Weierstraß is often stated in two separate forms. The first is that

if $[a,b]$ is a compact interval in $\mathbb{R}$ and $f:[a,b]\longrightarrow\mathbb{R}$ is continuous, then f is uniformly approximable on $[a,b]$ by polynomials; for any $\epsilon>0$, there is a polynomial p such that

$$(\forall t\in[a,b])\quad |p(t)-f(t)|<\epsilon\,.$$

Alternatively, for any continuous function $g:[-\pi,\pi]\longrightarrow\mathbb{C}$ such that $g(-\pi)=g(\pi)$, there is a "trigonometrical polynomial" $q(t)$ such that $(\forall t\in[-\pi,\pi])\ |q(t)-f(t)|<\epsilon$. By a trigonometrical polynomial, I mean a function of the form

$$\begin{aligned} &a_0+a_1\cos t+a_2\cos(2t)+\cdots+a_k\cos(kt)+\\ &\qquad+b_1\sin t+b_2\sin(2t)+\cdots+b_k\sin(kt) \end{aligned} \tag{19}$$

for some complex constants $a_0,a_1,\ldots,a_k,b_0,b_1,\ldots,b_k$ and some integer $k\in\mathbb{N}$.

This theorem, in either form, was originally proved by Weierstraß around 1885, and many proofs have been given by later mathematicians. There are generalizations to higher dimensions, and obvious questions, such as "for a given ϵ and f, what determines how high the degree of p must be, or what the coefficients must look like?", that have been much studied. But M. H. Stone, in 1948, pointed out that the pure approximation theorem, as distinct from questions of closeness of approximation, can be proved rather straightforwardly in a very general context. For many purposes his version is the most convenient, and it has become one of the most often cited of all theorems of abstract analysis. There are cleverer proofs than the one I shall give (which is Stone's in all essentials), but they assume much more. We need some preparatory ideas, of which the first is an often useful elementary result.

Theorem 6.1. (Dini's theorem.) *Let Ω be a compact topological space, and (f_n) a sequence in $C(\Omega;\mathbb{R})$. Suppose that f_n tends pointwise* monotonically *to $f\in C(\Omega;\mathbb{R})$. Then $f_n\to f$* uniformly.

The limit must be known to be continuous and the convergence must be monotonic.

Proof. By considering $|f_n-f|$ instead of f_n, we may as well assume that $f_n\downarrow 0$. Take any $\epsilon>0$, and define, for any $n\in\mathbb{N}$, $U(n,\epsilon):=\{x\in\Omega:\ f_n(x)<\epsilon\}$. This is open in Ω, as f_n is continuous. For any $x\in\Omega$, there is some n such that $f_n(x)<\epsilon$, so $\Omega=\bigcup_{n=1}^{\infty}U(n,\epsilon)$. As Ω is compact, there is a subcover consisting of finitely many of the sets $U(n,\epsilon)$; let N be the largest index appearing in this subcover, so $\Omega=\bigcup_{n=1}^{N}U(n,\epsilon)$.

Now, if $m \geq N$, and x is any point of Ω, there is some $n \leq N$ for which $x \in U(n, \epsilon)$, and so $0 \leq f_m(x) \leq f_N(x) \leq f_n(x) < \epsilon$. That is, $m \geq N \Longrightarrow (\forall x \in \Omega)\ 0 \leq f_m(x) < \epsilon$, or $f_n \to 0$ uniformly on Ω. $\square$

Lemma 6.2. *The function* $t \mapsto |t| : [-1,1] \longrightarrow \mathbb{R}$ *may be uniformly approximated by even polynomials with zero constant term.*

Proof. Define a sequence (p_n) of polynomials by induction, as follows.

$$p_0(t) = 0\,, \quad p_{n+1}(t) = p_n(t) + \tfrac{1}{2}(t^2 - (p_n(t))^2)\,. \tag{20}$$

It follows that each $p_n(t)$ is an even polynomial (only even powers of t appear) with zero constant term: $p_1(t) = \frac{1}{2}t^2$, $p_2(t) = t^2 - \frac{1}{8}t^4$, and so on.

Suppose that $0 \leq p_n(t) \leq |t|$ for $|t| \leq 1$ (which is certainly true for $n = 0$). Then (20) shows that $p_{n+1}(t) \geq p_n(t)$ for $|t| \leq 1$. Furthermore, again for $|t| \leq 1$,

$$\begin{aligned} p_{n+1}(t) - p_n(t) &= \tfrac{1}{2}(|t| + p_n(t))(|t| - p_n(t)) \\ &\leq |t|(|t| - p_n(t)) \leq |t| - p_n(t)\,, \end{aligned} \tag{21}$$

and so $p_{n+1}(t) \leq |t|$ for $|t| \leq 1$ too. By induction, then, for each $t \in [-1,1]$ and $n \geq 0$,

$$0 \leq p_n(t) \leq p_{n+1}(t) \leq |t|\,.$$

Hence, for each individual $t \in [-1,1]$, $p_n(t)$ increases monotonically to a limit $p(t) \leq |t|$.

Suppose that $|t| - p(t) = \kappa > 0$. Then by (21), for any n, $|t| - p_n(t) \geq \kappa$ and

$$p_{n+1}(t) - p_n(t) \geq \tfrac{1}{2}\kappa|t|\,,$$

which is absurd for $t \neq 0$, as $p_{n+1}(t) - p_n(t) \to 0$. Hence, $p(t) = |t|$, and $p_n(t) \uparrow |t|$ for $0 \leq t \leq 1$. As each p_n is even, the same is true for $-1 \leq t \leq 1$. The stated result now follows by 6.1. $\square$

Remark 6.3. There is an alternative proof of the above Lemma if one is willing to assume the binomial theorem for fractional exponent, specifically for exponent $\frac{1}{2}$, viz. that

$$(1+x)^{1/2} = 1 + \tfrac{1}{2}x + \frac{\frac{1}{2}(\frac{1}{2}-1)}{2!}x^2 + \frac{\frac{1}{2}(\frac{1}{2}-1)(\frac{1}{2}-2)}{3!}x^3 + \cdots$$

for $-1 < x < 1$. This is not, however, an entirely simple result. (It is even true for $x = \pm 1$, but that is still more difficult.) The series is uniformly convergent for $|x| < k$, if $0 < k < 1$.

Assuming these facts, and given $\epsilon \in (0,1)$, we may expand $\left(1 + \dfrac{y - \frac{1}{2}}{\epsilon^2 + \frac{1}{2}}\right)^{1/2}$ by the binomial series. It converges uniformly for $\dfrac{-1}{1+2\epsilon^2} \leq \dfrac{y - \frac{1}{2}}{\epsilon^2 + \frac{1}{2}} \leq \dfrac{1}{1+2\epsilon^2}$, i.e. for $0 \leq y \leq 1$. Take $y = t^2$, and then

$$(t^2 + \epsilon^2)^{1/2} = (\tfrac{1}{2} + \epsilon^2)^{1/2}\left(1 + \frac{t^2 - \frac{1}{2}}{\epsilon^2 + \frac{1}{2}}\right)^{1/2}$$

may be approximated uniformly in $[-1,1]$ by polynomials in t^2. In particular, there will be a polynomial $p_1(t)$ such that $\left|(t^2+\epsilon^2)^{1/2}-p_1(t)\right|<\epsilon$ for $-1\le t\le 1$. Then $|p_1(0)|<2\epsilon$. Set $p_2(t):=p_1(t)-p_1(0)$, an even polynomial with zero constant term. For $t\in[-1,1]$,

$$\begin{aligned}||t|-p_2(t)| &\le \left|(t^2+\epsilon^2)^{1/2}-p_1(t)\right|+|p_1(0)|+(t^2+\epsilon^2)^{1/2}-|t| \\ &< \epsilon+2\epsilon+\frac{(t^2+\epsilon^2)-|t|^2}{(t^2+\epsilon^2)^{1/2}+|t|}\le \epsilon+2\epsilon+\epsilon=4\epsilon\,.\end{aligned}$$

So p_2 will approximate $|t|$ uniformly on $[-1,1]$ within 4ϵ, which is good enough.

The idea of Stone's argument is that $C(\Omega;\mathbb{R})$ is a *lattice* as well as having the other structures we mentioned above.

Definition 6.4. Let $f,g\in C(\Omega;\mathbb{R})$. The *join* of f and g is $f\vee g\in C(\Omega;\mathbb{R})$, defined by

$$(\forall x\in\Omega)\quad (f\vee g)(x)=\max(f(x),g(x))\,,$$

and the *meet* of f and g is $f\wedge g\in C(\Omega;\mathbb{R})$, defined by

$$(\forall x\in\Omega)\quad (f\wedge g)(x)=\min(f(x),g(x))\,.$$

These are of course the lattice-theoretic names and notations for operations that are otherwise familiar. They are clearly commutative and associative. To complete the definition of a lattice, one would need to know that $\vee$ and $\wedge$ are *idempotent*:

$$(\forall f\in C(\Omega;\mathbb{R}))\ \ f\vee f=f\ \ \&\ \ f\wedge f=f\,,$$

and that they satisfy the *absorption identities*: $f\wedge(f\vee g)=f=f\vee(f\wedge g)$. These properties are obviously satisfied, so that $C(\Omega;\mathbb{R})$ is indeed a lattice in the algebraic sense.

Definition 6.5. A subset B of $C(\Omega;\mathbb{R})$ is a *sublattice* if, for any $f,g\in B$, both $f\vee g\in B$ and $f\wedge g\in B$.

We shall not need to go into the algebraic theory of lattices, and indeed we could describe $\vee$, $\wedge$ simply as the maximum and minimum, or as the *upper envelope* and *lower envelope*. [In principle, one might define the upper or lower envelope of any class of functions in $C(\Omega;\mathbb{R})$:

$$\left(\bigvee\nolimits_{\beta\in B} f_\beta\right)(x):=\sup\{f_\beta(x):\beta\in B\}\,,\qquad \left(\bigwedge\nolimits_{\beta\in B} f_\beta\right)(x):=\sup\{f_\beta(x):\beta\in B\}\,,$$

but the results of these operations need not be continuous, i.e. in $C(\Omega;\mathbb{R})$, unless B is finite, in which case they can be expressed by repeated application of $\vee$ and $\wedge$.]

Lemma 6.6. For any $f,g\in C(\Omega;\mathbb{R})$,

$$\begin{aligned}\tfrac12(f+g)+\tfrac12|f-g| &= f\vee g\,,\\ \tfrac12(f+g)-\tfrac12|f-g| &= f\wedge g\,. \text{ *******}\end{aligned}$$

□

Definition 6.7. Let $A\subseteq C(\Omega;\mathbb{K})$. A is a *subalgebra* of $C(\Omega;\mathbb{K})$ if it is a vector subspace over $\mathbb{K}$ and, for all $f,g\in A$, their product fg also belongs to A. [We do not require A to contain the constant functions, and therefore it need not have an "identity" element.]

A subset S of $C(\Omega;\mathbb{R})$ is said to *separate points*, or to be *separating*, if, for any $x,y\in\Omega$ such that $x\ne y$, there is some $f\in S$ such that $f(x)\ne f(y)$.

Lemma 6.8. *Let A be a subalgebra of $C(\Omega;\mathbb{R})$. Then $\mathrm{cl}(A)$, its closure with respect to the norm in $C(\Omega;\mathbb{R})$ (see* 4.3*), is also a subalgebra.*

Proof. Suppose that $f, g \in \mathrm{cl}(A)$ and $\lambda \in \mathbb{R}$. There are sequences (f_n) and (g_n) such that, for each $n \in \mathbb{N}$, $f_n, g_n \in A$ and $f_n \to f$, $g_n \to g$; that is, $\|f_n - f\| \to 0$, $\|g_n - g\| \to 0$. We may omit some terms at the beginning, and renumber, so as to ensure that $\|g_n - g\| \le 1$ for all n, and then, for all n,

$$\|g_n\| \le \|g\| + \|g_n - g\| \le K \coloneqq \|g\| + 1 .$$

However, $\lambda f_n, f_n g_n, f_n + g_n \in A$ for each n (as A is a subalgebra), and

$$\begin{aligned}
\|\lambda f_n - \lambda f\| &= |\lambda| \|f_n - f\| \to 0 , \\
\|f_n g_n - fg\| &= \|(f_n - f)g_n + f(g_n - g)\| \\
&\le \|f_n - f\|\|g_n\| + \|f\|\|g_n - g\| \\
&\le K\|f_n - f\| + \|f\|\|g_n - g\| \to 0 , \\
\|(f_n + g_n) - (f + g)\| &= \|(f_n - f) + (g_n - g)\| \\
&\le \|f_n - f\| + \|g_n - g\| \to 0 .
\end{aligned}$$

Thus, λf, $f + g, fg$ are limits of sequences in A, namely of $(\lambda f_n), (f_n + g_n), (f_n g_n)$. □

Lemma 6.9. *If A is a subalgebra of $C(\Omega;\mathbb{R})$ and $f \in A$, then $|f| \in \mathrm{cl}\, A$.*

[Here $|f|$ denotes the function in $C(\Omega;\mathbb{R})$ whose value at $x \in \Omega$ is $|f(x)|$.]

Proof. If necessary, multiply f by a non-zero scalar to ensure $\|f\| \le 1$. Take any $n \in \mathbb{N}$; by 6.2, there is an even polynomial q_n with zero constant term such that $|q_n(t) - |t|| < 2^{-n}$ for $-1 \le t \le 1$, and so, for any $x \in \Omega$, $|q_n(f(x)) - |f(x)|| < 2^{-n}$. Thus

$$\sup\{|q_n(f(x)) - |f(x)||\} < 2^{-n} , \tag{22}$$

the strict inequality applying because of 3.15*(a)* — the supremum is $|q_n(f(x_0)) - |f(x_0)||$ for some suitable $x_0 \in \Omega$.

$q_n(f)$ may be interpreted in A in an obvious way. If the polynomial q_n were given by

$$q_n(t) = \alpha_1 t + \alpha_2 t^2 + \cdots + \alpha_k t^k$$

(in fact the coefficients of odd powers are all 0 in this case), then

$$q_n(f) \coloneqq \alpha_1 f + \alpha_2 f^2 + \cdots + \alpha_k f^k \in A .$$

Thus, $(q_n(f))(x) = q_n(f(x))$ for each $x \in \Omega$. Thus (22) tells us

$$\|q_n(f) - |f|\| < 2^{-n} ,$$

which shows that $q_n(f) \to |f|$ and completes the proof. □

Corollary 6.10. *Let A be a subalgebra of $C(\Omega;\mathbb{R})$. If $f, g \in A$, then $f \vee g, f \wedge g \in \mathrm{cl}\, A$. In particular, if B is a closed subalgebra and $f, g \in B$, then $f \vee g, f \wedge g \in B$; that is, any closed subalgebra B is also a sublattice of $C(\Omega;\mathbb{R})$.*

Proof. By 6.6, 6.8, and 6.9. □

Lemma 6.11. *Let B be a sublattice of $C(\Omega;\mathbb{R})$. Suppose that $h \in C(\Omega;\mathbb{R})$ and that, for any two points x, y of Ω and for any $\epsilon > 0$, there exists some $g \in B$ such that*

$$|h(x) - g(x)| < \epsilon, \quad |h(y) - g(y)| < \epsilon. \tag{23}$$

Then $h \in \mathrm{cl}\, B$.

[In other words: if h may be approximated at *pairs of points* by elements of the sublattice, then it may be approximated *uniformly* on the whole of Ω. This is Stone's crucial observation.]

Proof. Fix $\epsilon > 0$, and, for given $x, y \in \Omega$, let $g_{xy} \in B$ be such that $|h(x) - g_{xy}(x)| < \epsilon$ and $|h(y) - g_{xy}(y)| < \epsilon$. Let

$$U(x,y) := \{u \in \Omega : g_{xy}(u) < h(u) + \epsilon\};$$

this is an open set in Ω containing both x and y. Hence, holding x fixed, $\{U(x,y) : y \in \Omega\}$ is an open cover of Ω, and there is a finite subcover $\{U(x,y_1), \dots, U(x,y_m)\}$ say. Let

$$g_x := g_{x,y_1} \wedge g_{x,y_2} \wedge \cdots \wedge g_{x,y_m} \in B.$$

For any $y \in \Omega$, there is some k such that $y \in U(x, y_k)$, and then

$$g_x(y) \le g_{x,y_k}(y) < h(y) + \epsilon. \tag{24}$$

On the other hand, for any $x \in \Omega$ the choice of g_{x,y_k} for each k gives

$$g_x(x) = \min(g_{x,y_1}(x), g_{x,y_2}(x), \dots, g_{x,y_m}(x)) > h(x) - \epsilon,$$

as each g_{x,y_k} was chosen so that $g_{x,y_k}(x) > h(x) - \epsilon$, and so

$$x \in V(x) := \{v \in \Omega : g_x(v) > h(v) - \epsilon\}.$$

But $\{V(x) : x \in \Omega\}$ is an open covering of Ω, so it has a finite subcovering $\{V(x_i) : 1 \le i \le n\}$. Let

$$g := g_{x_1} \vee g_{x_2} \vee \cdots \vee g_{x_n} \in B,$$

and then $g(x) > h(x) - \epsilon$ for all $x \in \Omega$. But also $g(x) < h(x) + \epsilon$ for all x, because this inequality is true for each g_{x_i} individually, by (24). Hence $|h(x) - g(x)| < \epsilon$ for all $x \in \Omega$, or $\|h - g\| < \epsilon$, where $g \in B$.

Since such a g may be constructed for any $\epsilon > 0$, it follows that $h \in \mathrm{cl}\, B$. □

Theorem 6.12. (M. H. Stone.) *Let Ω be a compact topological space, and A a separating subalgebra of $C(\Omega;\mathbb{R})$. Either A is dense in $C(\Omega;\mathbb{R})$, $\mathrm{cl}(A) = C(\Omega;\mathbb{R})$, or else there is a point $\omega \in \Omega$ such that $\mathrm{cl}(A)$ consists of all the functions in $C(\Omega;\mathbb{R})$ that vanish at ω.*

Proof. Suppose first that there is no point at which all the functions in A vanish. Then, if $x \ne y$ in Ω, there are functions $f_1, f_2, f_3 \in A$ such that $f_3(x) \ne f_3(y)$ (A separates points) and $f_1(x) \ne 0 \ne f_2(y)$.

There exists $f \in A$ such that $0 \neq f(x) \neq f(y) \neq 0$. Indeed, f_3 may be taken for f, unless either $f_3(x) = 0$ or $f_3(y) = 0$. If $f_3(x) = 0$ and consequently $f_3(y) \neq 0$, let $\alpha > 0$ be such that $\alpha|f_1(x) - f_1(y)| < |f_3(y)| > \alpha|f_1(y)|$, and take $f := \alpha f_1 + f_3 \in A$. It follows that $f(x) = \alpha f_1(x) \neq 0$, $f(y) = \alpha f_1(y) + f_3(y) \neq 0$, and $f(x) \neq f(y)$. If $f_3(y) = 0$, argue symmetrically.

Suppose $a, b \in \mathbb{R}$. There exists $g \in A$ such that $g(x) = a$ and $g(y) = b$. To establish this, try setting $g := \lambda f + \mu f^2 \in A$; it will suffice if $\lambda, \mu \in \mathbb{R}$ satisfy

$$\lambda(f(x)) + \mu(f(x))^2 = a\,, \quad \lambda(f(y)) + \mu(f(y))^2 = b\,.$$

As a system of equations in λ and μ, these have a unique solution, since the determinant of the coefficient matrix is $f(x)f(y)(f(x) - f(y)) \neq 0$.

By 6.10, $\mathrm{cl}\,A$ is a sublattice of $C(\Omega;\mathbb{R})$. From this and the previous paragraph, the hypotheses of 6.11 hold with $B := \mathrm{cl}\,A$; indeed, (23) holds with $h(x) = g(x)$ and $h(y) = g(y)$. So any member of $C(\Omega;\mathbb{R})$ is in $\mathrm{cl}\,B = \mathrm{cl}(\mathrm{cl}\,A) = \mathrm{cl}\,A$; A is dense.

This leaves the case where there is some point $\omega \in \Omega$ such that every function in A vanishes at ω. Let A_1 be the set of functions of the form $\phi + c$ for $\phi \in A$ and $c \in \mathbb{R}$, by which I mean the function defined by $(\phi + c)(x) := \phi(x) + c$ for each $x \in \Omega$. It is easily checked that A_1 is also a subalgebra of $C(\Omega;\mathbb{R})$; it obviously includes A, and so separates points; and there is no point of Ω at which all its members vanish. From the first case, above, $\mathrm{cl}(A_1) = C(\Omega;\mathbb{R})$. But it is clear that $\mathrm{cl}(A)$ can contain only functions that vanish at ω.

Suppose $\psi \in C(\Omega;\mathbb{R})$, $\psi(\omega) = 0$, and $\epsilon > 0$. There exist $c \in \mathbb{R}$ and $\phi \in A$ such that $\|\psi - (\phi + c)\| < \frac{1}{2}\epsilon$. But then $\|c\| = |c| = |\psi(\omega) - (\phi(\omega) + c)| < \frac{1}{2}\epsilon$, and so

$$\|\phi - \psi\| \leq \|\psi - (\phi + c)\| + \|c\| < \tfrac{1}{2}\epsilon + \tfrac{1}{2}\epsilon = \epsilon\,.$$

Therefore, any continuous function ψ vanishing at ω belongs to the closure of A.

Corollary 6.13. *Let Ω be a compact topological space and let D be a separating subset of $C(\Omega;\mathbb{R})$. If there is no point at which all the functions in D vanish, then any continuous function $\Omega \longrightarrow \mathbb{R}$ may be uniformly approximated on Ω by polynomials with zero constant term in the functions of D. If all the functions in D vanish at some point $\omega \in \Omega$, then any continuous function vanishing at ω may be uniformly approximated by polynomials with zero constant term in the functions of D, and any continuous function may be uniformly approximated by polynomials in the functions of D that may have non-zero constant term.*

Proof. The set of polynomials described is in each case a separating subalgebra. So this is just a rephrasing of 6.12. □

This Corollary instantly yields Weierstrass's theorems. For the first theorem, take $\Omega := [a, b]$, and let D be the singleton $\{\chi\}$, where $\chi(t) = t$ for all $t \in [a, b]$.

Now take $\Omega := \mathbb{T}$, the unit circle in the complex plane $\{z \in \mathbb{C} : |z| = 1\}$, and let $E : [-\pi, \pi] \longrightarrow \mathbb{T} : t \mapsto \exp(it)$. A continuous function $g : [-\pi, \pi] \longrightarrow \mathbb{R}$ such that $g(-\pi) = g(\pi)$, such as was considered in the second version of Weierstraß's theorem, may be "factored through" E; that is, there is a unique continuous function $f : \mathbb{T} \longrightarrow \mathbb{R}$ such that $f \circ E = g$. In this way, we reduce the problem to one concerning $C(\mathbb{T};\mathbb{R})$. The function $\cos(nt)$ on $[-\pi, \pi]$ corresponds to $\Re(z^n)$ on $\mathbb{T}$: $\Re(E(t)^n) = \cos(nt)$, and similarly $\sin(nt)$ corresponds to $\Im(z^n)$. In any case, the trigonometrical polynomials (19), regarded in this way as continuous functions on $\mathbb{T}$, form a subalgebra of $C(\mathbb{T};\mathbb{R})$. It separates points; indeed, $\Re(z)$ and $\Im(z)$ are sufficient for this. So, by Stone's theorem, any continuous function $T \longrightarrow \mathbb{R}$ may be uniformly approximated by real linear combinations of the functions $\Re(z^n)$ and $\Im(z^n)$ for $n = 0, 1, 2, \dots$, or, equivalently. any continuous function $g : [-\pi, \pi] \longrightarrow \mathbb{R}$ such that $g(-\pi) = g(\pi)$ may be uniformly approximated by trigonometrical polynomials.

Remark 6.14. The remarkable things about Stone's theorem are its generality and its elementary character. They are related, of course, for a result so general is unlikely to involve "hard analysis". But, as a consequence, it cannot answer any of the hard questions that arise in specific instances, which earlier proofs of Weierstrass's theorems, for instance, may give some information on. As an example of what I mean, consider the function $\cos x$; we know there is some polynomial $q_\epsilon(x)$ in x such that

$$(\forall x \in [-\pi, \pi]) \quad |q_\epsilon(x) - \cos x| < \epsilon\,,$$

but it is reasonable to ask, for a given ϵ, how small the degree of q_ϵ may be, or how large its coefficients may have to be. Stone's theorem can give no answer, as is inevitable in a result which applies to any continuous function whatever. More specialized arguments (for instance from the cosine series) can be used; but they will involve "hard analysis" at some point.

The other obvious defect of Stone's theorem is that, in using the lattice structure of a subalgebra, the proof is restricted to the real case. Not just the proof fails, however; the result itself is untrue for complex algebras of functions.

Let $\Omega := \mathbb{T}$, and consider $C(\Omega; \mathbb{C})$. It has a subalgebra A consisting of the restrictions to $\mathbb{T}$ of complex polynomials $a_0 + a_1 z + a_2 z^2 + \cdots + a_k z^k$, with $a_0, a_1, \dots, a_k \in \mathbb{C}$. The mapping $f : z \mapsto \overline{z}$ is in $C(\Omega; \mathbb{C})$. However, consider the complex contour integrals:

$$\int_{\mathbb{T}} f(z)\,dz = \int_{\mathbb{T}} \overline{z}\,dz = \int_{\mathbb{T}} \frac{dz}{z} = 2\pi i\,,$$

whilst, for any polynomial $p(z)$, $\displaystyle\int_{\mathbb{T}} p(z)\,dz = 0$ (because p is defined and holomorphic on the whole of $\mathbb{C}$). It is therefore impossible to approximate f uniformly on $\mathbb{T}$ by complex polynomials. If $|f(z) - p(z)| < \epsilon$ for all $z \in \mathbb{T}$, then, by the "fundamental estimate",

$$2\pi = |2\pi i| = \left| \int_{\mathbb{T}} f(z)\,dz - \int_{\mathbb{T}} p(z)\,dz \right| \leq \int_{\mathbb{T}} |f(z) - p(z)||dz| \leq 2\pi\epsilon\,,$$

which is absurd if $\epsilon < 1$. To obtain an analogue of Stone's theorem in the complex case, an extra hypothesis is necessary.

Definition 6.15. Let A be a complex linear subspace of $C(\Omega; \mathbb{C})$, where Ω is a compact topological space. A is *self-conjugate* if, for any $f \in A$, the function $\overline{f} : \Omega \longrightarrow \mathbb{C}$ defined by

$$(\forall x \in \Omega) \quad \overline{f}(x) := \overline{f(x)}$$

also belongs to A.

Theorem 6.16. (Stone's theorem, complex version). *Let Ω be a compact topological space, and let A be a separating self-conjugate complex subalgebra of $C(\Omega; \mathbb{C})$. Then either the closure of A in the supremum norm is $C(\Omega; \mathbb{C})$, or there is a point $\omega \in \Omega$ such that the closure of A consists of all the continuous functions $f : \Omega \longrightarrow \mathbb{C}$ for which $f(\omega) = 0$.*

Proof. If $f \in A$, then $\Re(f) = \frac{1}{2}(f + \overline{f}) \in A$, and $\Im(f) = -\frac{1}{2}i(f - \overline{f}) \in A$. The real-valued functions in A form a *real* subalgebra A_0 of A, which must separate points: if $f(x) \neq f(y)$, where $f \in A$, then either $\Re(f)(x) \neq \Re(f)(y)$ or $\Im(f)(x) \neq \Im(f)(y)$.

Suppose that, for any $x \in \Omega$, there is some $f \in A$ with $f(x) \neq 0$; then either $\Im(f)(x) \neq 0$ or $\Re(f)(x) \neq 0$. By 6.12, the closure of A_0 in the uniform norm is the whole of $C(\Omega; \mathbb{R})$. Take any $g \in A$; then, for any $\epsilon > 0$, there exist $g_1, g_2 \in A_0$ such that

$$\|\Re(g) - g_1\| < \tfrac{1}{2}\epsilon, \quad \|\Im(g) - g_2\| < \tfrac{1}{2}\epsilon.$$

Consequently $\|g - (g_1 + ig_2)\| < \epsilon$, and $g_1 + ig_2 \in A$.

Finally, suppose that, for some $\omega \in \Omega$, $f(\omega) = 0$ for all $f \in A$. In this case the closure of A_0 consists of all real-valued continuous functions on Ω that vanish at ω, and, as before, the closure of A must consist of all $f \in C(\Omega; \mathbb{C})$ such that $f(\omega) = 0$. □

Corollary 6.17. *Let B be a separating subset of $C(\Omega; \mathbb{C})$. Any $f \in C(\Omega; \mathbb{C})$ may be uniformly approximated by complex polynomials in the members of B and their conjugates. If, for every $x \in \Omega$, there is some $h \in B$ such that $h(x) \neq 0$, then the polynomials in question may always be chosen to have zero constant term.* □

There is a further extension of Stone's theorem. As it stands it cannot apply to continuous functions on $\mathbb{R}$, which is not compact. There is a trifling modification of the theorem that does work in such cases, but it requires a brief introduction.

Definition 6.18. Let Ω be a topological space. Define Ω^* to be the set $\Omega \cup \{*\}$, where $* \notin \Omega$, with a topology whose open sets are the open sets of Ω and the complements (in $\Omega \cup \{*\}$) of those sets that are both closed and compact in Ω.

It is an easy exercise to show that Ω^* is a topological space, and that it is compact. It is called the *one-point compactification* or *Alexandrov compactification* of Ω, since it is obtained by adding a single point "$*$" to Ω whilst keeping the same subsets of Ω as open. [There is a slight logical difficulty in demanding that $* \notin \Omega$. But, for instance, there must be a subset of Ω which is not also an element of Ω, since by Cantor's theorem $\mathcal{P}(\Omega)$ has more elements than Ω; and we could call it $*$. It is unimportant what $*$ is, provided it is "new".]

Definition 6.19. A continuous function $f : \Omega \longrightarrow \mathbb{K}$ is said to *vanish at infinity* if, for every $\epsilon > 0$, the set $\{x \in \Omega : |f(x)| \geq \epsilon\}$ is compact. Such functions form an algebra over $\mathbb{K}$.

Lemma 6.20. *If $f : \Omega \longrightarrow \mathbb{K}$ is continuous, then it vanishes at infinity if and only if there is a continuous function $f^* : \Omega^* \longrightarrow \mathbb{K}$ such that $f(*) = 0$ and $f^*|\Omega = f$. Furthermore, f^* is unique, given f.* □

It follows that such functions f are bounded, form an algebra $C_0(\Omega; \mathbb{K})$ under pointwise operations, and can be normed by a supremum norm as usual.

Proposition 6.21. *Let A be a subalgebra of $C_0(\Omega; \mathbb{K})$ which separates points; in the case $\mathbb{K} = \mathbb{C}$, is self-conjugate; and, for any $x \in \Omega$, contains an element that does not vanish at x. Then any $f \in C_0(\Omega; \mathbb{K})$ may be uniformly approximated on Ω by elements of A.*

Proof. Extending each function in A to Ω^*, we obtain a subalgebra of $C(\Omega^*; \mathbb{K})$ which separates points — and all its members vanish at $*$. The result follows from 6.12. □

[The hypotheses can, in fact, only be satisfied if Ω is locally compact.***]

Example 6.22. Let $\Omega := [0, \infty)$. The function $\exp(-\sqrt{x})$ vanishes at infinity on $\mathbb{R}$. A may be the set of polynomials in $\exp(-x)$, which itself is everywhere non-zero and separates points. Thus, $e^{-\sqrt{x}}$ can be uniformly approximated on $[0, \infty)$ by polynomials in e^{-x}.

The main limitation of Stone's theorem is its restriction to *algebras* of functions. As an example of the sort of question it does *not* answer, consider this. Suppose we have a sequence $\alpha(1), \alpha(2), \alpha(3), \ldots$ of distinct positive numbers. Do the functions $t^{\alpha(1)}, t^{\alpha(2)}, t^{\alpha(3)}, \ldots$ span a dense vector subspace of $C([1,2];\mathbb{R})$? Weierstraß's theorem (for polynomials) says the answer is positive if $\alpha(1), \alpha(2), \alpha(3), \ldots$ is an arithmetic progression (or a rearrangement of one), but, for instance, is it positive if $\alpha(k) := 2^k$ for each k? [The answer to this question is known; it is a theorem of Szász from about 1918. But the proof is not easy. There is another proof by Paley and Wiener in their book. Problems of this kind are not entirely without practical significance.]

§7. Equivalents of the axiom of choice.

As in 441, I do not wish to spend much time on the axiom of choice as such, because it is presumably fully treated in logic courses. But in 441 I made use only of the "pure" axiom, which happened to be convenient there, and in other contexts there are other statements that are more useful and in a sense equivalent. The "pure" Axiom of Choice is the assertion that, for any set $\mathcal{C}$ whose members are non-empty sets, there is a function $f : \mathcal{C} \longrightarrow \bigcup_{C \in \mathcal{C}} C$ such that $f(C) \in C$ for each $C \in \mathcal{C}$. (In other words, one may "choose" one element, by the "choice function" f, from each of the sets $C \in \mathcal{C}$. Notice that C in $f(C)$ is simply the value of the variable; the arguments of f are members of $\mathcal{C}$.)

The prominence of the Axiom of Choice is due, historically speaking, to its asserting the existence of a "choice function" even in cases where it is inconceivable to construct one in any "practical" way. At the merely intuitive level, I find this unexceptionable, and perhaps no-one would have worried very much about it had not Russell constructed his paradox and thereby demonstrated the pitfalls of naive set theory. We do know, now, that the axiom is consistent (as is its negation) with the other customarily accepted axioms of set theory. However, there is a large number of statements which were proved quite early on to be equivalent to the Axiom of Choice, in the sense that they may all be derived from each other via standard set theory. I shall not give proofs; you can find a concise discussion in Kelley's book *General Topology* and in many other places, and, in truth, we do not need more than a smattering of information.

A. The *Multiplicative Axiom* is the assertion that, if $\{C_\alpha : \alpha \in A\}$ is an indexed class of non-empty sets, the product $\prod_{\alpha \in A} C_\alpha$ (or $\times_{\alpha \in A} C_\alpha$) is also non-empty.

The point at issue here is the definition of the Cartesian product of a general indexed family, not just of finitely many sets $C_1, C_2, \ldots, C_k$. The usual convention is that $\prod_{\alpha \in A} C_\alpha$ is *by definition* the set of functions $A \longrightarrow \bigcup_{\alpha \in A} C$ for which $f(\alpha) \in C_\alpha$ for each $\alpha \in A$. (Thus, $\prod_{\alpha = 1,2} C_\alpha$ is the set of functions $f : \{1,2\} \longrightarrow C_1 \cup C_2$ for which $f(1) \in C_1$ and $f(2) \in C_2$. It is clear that this is, as it were, "functionally equivalent" to the definition of $C_1 \times C_2$ as a set of ordered pairs, although it is *not* the same.) This being so, the Multiplicative Axiom, as I have stated it, is almost a rephrasing of the Axiom of Choice, as I have stated it, and they are almost trivially equivalent. [What is the difference between them?]

B. *Tikhonov's theorem* is the assertion that, if $\{\Omega_\alpha : \alpha \in A\}$ is an indexed class of topological spaces and each Ω_α is compact, then the product topological space $\prod_{\alpha\in A}\Omega_\alpha$ is also compact. A full statement of this theorem would demand a more detailed discussion of the product topology than I wish to give here (see 10.24).

The theorem is equivalent to the Axiom of Choice. (It is quite easy, with a little knowledge of topology, to see that it implies the Multiplicative Axiom; the converse proof is non-trivial.)

C. The *Well-Ordering Principle* (also known as *Zermelo's Axiom*, or indeed *Zermelo's Theorem*) states that any set may be well-ordered.

A *partial order* on a set S is a binary relation $\leq$ on S which is reflexive ($a \leq a$) and transitive ($a \leq b$ & $b \leq c \Longrightarrow a \leq c$), such that $a \leq b$ & $b \leq a \Longrightarrow a = b$. (The last condition excludes trivial examples.) A *well-ordering* on S is a partial order such that every subset of S has a least element.

The equivalence of the Axiom of Choice and the Well-Ordering Principle is well-known; in essence it goes back to Zermelo (1908), and may be found in many places — a sketch is in Kelley's book, for instance. The WOP is interesting because it allows one to carry out arguments by so-called transfinite induction, which extends ordinary induction, and so raises all the problems of transfinite arithmetic.

D. *Zorn's Lemma* is perhaps the most generally useful of these statements.

Let $(S, \leq)$ be any partially ordered set. A *chain* in S is a subset T that is totally ordered by $\leq$ (that is, if $x, y \in T$, then either $x \leq y$ or $y \leq x$). An *upper bound* for T is an element $b \in S$ such that, for any $x \in T$, $x \leq b$. A *maximal element* of S is an element $m \in S$ such that $m \leq y \in S \Longrightarrow y = m$. (Notice that it need not be an upper bound for S. For instance, let S be $\{\{1\},\{1,2\},\{1,3\}\}$, with $\subseteq$ as $\leq$. Then $\{1,2\}$ and $\{1,3\}$ are both maximal in S, but neither is greater than the other.) The partial order $\leq$ on S is *inductive* if every chain in S has an upper bound. Zorn's Lemma then says that

an inductive partial order on a non-empty set has a maximal element.

The statement is superficially plausible. If there were no maximal element, then any $c \in S$ must allow a greater element, and one could construct a *sequence* of elements of S that is a chain and must have an upper bound. If a large enough chain exists, its upper bound should be maximal in S. This is where the Axiom of Choice is needed. Again, Zorn's Lemma is equivalent to the Axiom, and proofs may be found in many places.

It should be emphasized that S has to be a *set* for Zorn's Lemma to work. (Otherwise one can easily derive paradoxes such as Cantor's or Burali-Forti's.)

The usefulness of Zorn's Lemma, which may at first sight seem decidedly over-elaborate, is precisely that it can be applied in such diverse situations by suitable choices of the partially ordered set. There are at least three other common statements equivalent to the Axiom of Choice (the Hausdorff maximal principle, Kuratowski's maximal principle, and the minimal principle) that are very similar to Zorn's Lemma and in fact are special cases of it.

To convince you that Zorn's lemma is useful, here are four remarkable applications of what is sometimes flippantly called Zornification.

Lemma 7.1. *Let $\mathbb{K}$ be any field. Any vector space V over $\mathbb{K}$ has a basis (in the algebraic sense, i.e. a linearly independent spanning set). Indeed, any linearly independent subset of V is included in a basis.*

Proof. Suppose S is a linearly independent set in V (for instance the null set). Let $\mathcal{L}$ be the class of linearly independent sets in V that include S. Clearly $S \in \mathcal{L}$. Define the partial order in $\mathcal{L}$ by letting "$L_1 \le L_2$" mean $L_1 \subseteq L_2$.

Now suppose that $\mathcal{C}$ is a non-empty chain in $\mathcal{L}$. Define

$$M := \bigcup_{L \in \mathcal{C}} L.$$

We shall show that M is itself linearly independent. Suppose, then, that $v_1, v_2, \dots, v_n$ are distinct members of M and $\lambda_1 v_1 + \lambda_2 v_2 + \cdots + \lambda_n v_n = 0$, where $\lambda_1, \lambda_2, \dots, \lambda_n \in \mathbb{K}$.

For each k, $1 \le k \le n$, there is some $L_k \in \mathcal{C}$ such that $v_k \in L_k$. But, as $\mathcal{C}$ is a chain, we can arrange the members L_k in "ascending order"; after rearranging and renumbering, $L_1 \le L_2 \le \cdots \le L_n$. But this means that $v_1, v_2, \dots, v_n \in L_n$, and, as L_n is linearly independent, $\lambda_1 v_1 + \lambda_2 v_2 + \cdots + \lambda_n v_n = 0$ implies that $\lambda_1 = \lambda_2 = \cdots = \lambda_n = 0$. This proves that M is linearly independent, $M \in \mathcal{L}$, and M is obviously an upper bound for $\mathcal{C}$ in $\mathcal{L}$ (in fact the least upper bound); in particular, $S \subseteq M$.

By Zorn's Lemma, there is a maximal element B of $\mathcal{L}$, i.e. a maximal linearly independent subset of V. Suppose $x \in V$. Then $B \cup \{x\}$ cannot be linearly independent (or B would not be maximal); so there must exist distinct elements $w_1, w_2, \dots, w_k \in B \cup \{x\}$ and scalars $\mu_1, \mu_2, \dots, \mu_k \in \mathbb{K}$, not all zero, such that $\mu_1 w_1 + \mu_2 w_2 + \cdots + \mu_k w_k = 0$. One of the ws with a non-zero scalar coefficient must be x, since otherwise we have a non-trivial linear relation amongst the elements of B. Suppose, for instance, that it is w_k. Then

$$x = w_k = (-\mu_k^{-1}\mu_1)w_1 + \cdots + (-\mu_k^{-1}\mu_{k-1})w_k.$$

So x is a linear combination of the elements of B.

Hence, B is both linearly independent and spans V. □

It will be observed that the crucial point in showing that the partial order on $\mathcal{L}$ is inductive is that "linear independence of M" is equivalent to "linear independence of all finite subsets of M". There is yet another statement equivalent to the Axiom of Choice, the Teichmüller-Tukey Lemma, which refers explicitly to properties "of finite character" like this.

Lemma 7.2. *Let R be a ring with an identity e, and suppose that J is a proper left ideal in R. Then J is included in a maximal proper left ideal.*

Proof. Let $\mathcal{L}$ be the class of all proper left ideals that include J. Partially order $\mathcal{L}$ by inclusion: $J_1 \le J_2$ is to mean $J_1 \subseteq J_2$. If $\mathcal{C}$ is a non-empty chain in $\mathcal{L}$, then $\bigcup_{C \in \mathcal{C}} C$ is also an ideal in R. Indeed: if $x, y \in \bigcup_{C \in \mathcal{C}} C$, then there are $C_1, C_2 \in \mathcal{C}$ such that $x \in C_1$ and $y \in C_2$; but either $C_1 \subseteq C_2$ or $C_2 \subseteq C_1$; if, for example, $C_1 \subseteq C_2$, then $x, y \in C_2$, and so $x - y \in C_2 \subseteq \bigcup_{C \in \mathcal{C}} C$; if $r \in R$ and $x \in \bigcup_{C \in \mathcal{C}} C$, there is some $C \in \mathcal{C}$ such that $x \in C$, and $rx \in C \subseteq \bigcup_{C \in \mathcal{C}} C$. Of course $J \subseteq \bigcup_{C \in \mathcal{C}} C$. The one difficulty is in seeing that $\bigcup_{C \in \mathcal{C}} C$ is a *proper* left ideal. However, as each $C \in \mathcal{C}$ is a proper left ideal, $e \notin C$, and so $e \notin \bigcup_{C \in \mathcal{C}} C$.

Hence, the partial order on $\mathcal{L}$ is inductive, and by Zorn's Lemma it has a maximal element, which is precisely a maximal proper left ideal of R. □

It is usual to say "maximal left ideal" rather than "maximal proper left ideal", since, if R were allowed to be called an "ideal", it would itself be the only "maximal left ideal". The above result does have serious applications in analysis; but the next one is more purely algebraic (at least at the moment).

Definition 7.3. Suppose that R is a commutative ring with identity e. The *radical* $Z(R)$ of R is the set of all nilpotent elements of R:

$$Z(R) := \{x \in R : (\exists n \in \mathbb{N}) \;\; x^n = 0\}.$$

An ideal J of R is *prime* if a product ab (where $a, b \in R$) can belong to J only if at least one of the factors a, b belongs to J. [I do not claim there are any proper prime ideals.]

The radical $Z(R)$ is trivially an ideal in R, because of the commutativity of R. And J is a prime ideal if and only if the quotient ring R/J is an integral domain.

Lemma 7.4. *The radical of a commutative ring with identity is the intersection of all its prime ideals.*

Proof. If $x \in Z(R)$, then $x^n = 0$ for some $n \in \mathbb{N}$. Take any prime ideal J. Then $x^n \in J$, which clearly implies $x \in J$. Hence $Z(R) \subseteq J$, and $Z(R)$ is *included* in the intersection M of all the prime ideals, $Z(R) \subseteq M$.

Now, suppose that $x \notin Z(R)$. This means that all powers of x are non-zero.

Let $\mathcal{Q}$ be the class of all ideals of R which contain no power of x; partially order $\mathcal{Q}$ by inclusion. As in 7.2, the union of any chain in $\mathcal{Q}$ also belongs to $\mathcal{Q}$ and is, therefore, an upper bound for the chain in $\mathcal{Q}$. By Zorn's Lemma, there is a maximal element J in $\mathcal{Q}$. This means that any ideal that strictly includes J is not in $\mathcal{Q}$ and must contain a power of x.

Suppose $a \notin J$ and $b \notin J$. Then the ideal generated by a and J,

$$(a) + J = \{ra + y : r \in R \;\; \& \;\; y \in J\},$$

must contain $x^m = r_1 a + y_1$ for some $m \in \mathbb{N}$, $r_1 \in R$, $y_1 \in J$; and, similarly, for some $n \in \mathbb{N}$, $x^n = r_2 b + y_2$ for some $n \in \mathbb{N}$, $r_2 \in R$, $y_2 \in J$. Multiplying,

$$x^{m+n} = r_1 r_2 ab + r_1 a y_2 + r_2 b y_1 + y_1 y_2 \in (ab) + J.$$

However, $x^{m+n} \notin J$, and it must, consequently, be the case that $ab \notin J$. Hence,

$$a \notin J \;\; \& \;\; b \notin J \Longrightarrow ab \notin J,$$

which shows that J is prime.

Thus, if $x \notin Z(R)$, there is a prime ideal J with $x \notin J$, and $x \notin M$: $M \subseteq Z(M)$. ☐

The last example is definitely related to analysis.

Definition 7.5. Let Ω be a compact topological space, and let $\mathcal{F}$ be a non-empty subset of $C(\Omega; \mathbb{C})$. A *peak set* for $\mathcal{F}$ is a closed subset F of Ω with the property that, for every $f \in \mathcal{F}$, there exists a point $x \in F$ such that $|f(x)| = \sup\{|f(t)| : t \in \Omega\}$.

In view of 3.16*(a)*, this is equivalent to

$$\sup\{|f(x)| : x \in F\} = \sup\{|f(t)| : t \in \Omega\}.$$

F is a "peak set" in the sense that the moduli of the functions of $\mathcal{F}$ attain their suprema on F.

Lemma 7.6. *There exists a minimal peak set for $\mathcal{F}$.*

Proof. Let $\mathcal{L}$ be the class of peak sets. Certainly $\Omega \in \mathcal{L}$. Define $L_1 \leq L_2$, for $L_1, L_2 \in \mathcal{L}$, to mean $L_1 \supseteq L_2$ (the other way round from the previous examples). Let $\mathcal{C}$ be a chain in $\mathcal{L}$.

Take any $f \in \mathcal{F}$, and set $A_f := f^{-1}\{\sup\{|f(t)| : t \in \Omega\}\}$. By 3.5, A_f is closed in Ω. Thus, for any $K \in \mathcal{C}$, $A_f \cap K \neq \emptyset$ by the definition of a peak set. 3.16*(b)* implies that $A_f \cap (\bigcap_{K\in\mathcal{C}} K) \neq \emptyset$; and, as this holds for any $f \in \mathcal{F}$, $\bigcap_{K\in\mathcal{C}} K$ is also a peak set for $\mathcal{F}$. It is therefore an upper bound for $\mathcal{C}$ in $\mathcal{L}$. The result follows by Zorn's Lemma. $\square$

§8. Linear spaces.

As before, let $\mathbb{K}$ denote the field $\mathbb{R}$ or the field $\mathbb{C}$. Functional analysis began with Hilbert spaces — in fact with the space l^2 — and only later did Banach and his collaborators point out that many theorems could be proved more generally. The original Hilbert space proofs still sometimes appear in books that deal only with Hilbert spaces.

I have already assumed you know the axioms for a *vector space over* $\mathbb{K}$. Following the modern custom, I shall denote the "zero vector" in a vector space simply by 0, distinguishing it neither from the zero vectors in other vector spaces nor from the scalar "zero". In practice the context almost always indicates what is meant, and the device some authors have employed of writing "θ" for a zero *vector* has always seemed to me rather absurd — why θ rather than ω, say?

There are four elementary ways of obtaining "new vector spaces from old". I shall omit the details, which are rather lengthy, since the definition of a vector space is itself rather long.

1) A *vector subspace* of a vector space E is a subset F of E which contains the zero vector of E and has the property that, for any $x, y \in F$ and $\lambda, \mu \in \mathbb{K}$, $\lambda x + \mu y \in F$ too. It is then a vector space in its own right, with the same definitions of addition and scalar multiplication as for E.

2) If E_1, E_2 are vector spaces, the Cartesian product $E_1 \times E_2$ can be given the structure of a vector space by setting $\lambda(x_1, x_2) + \mu(y_1, y_2) := (\lambda x_1 + \mu y_1, \lambda x_2 + \mu y_2)$. The resulting vector space is the *direct sum* of E_1 and E_2, sometimes denoted $E_1 \oplus E_2$.

There are two ways of generalizing this idea to larger classes of vector spaces.

2a) Suppose that $\{E_\alpha : \alpha \in A\}$ is an indexed family of vector spaces over $\mathbb{K}$. Linear operations in the Cartesian product $\bigtimes_{\alpha\in A} E_\alpha$ (see §7) can be defined "componentwise" exactly as above, and this gives the "direct product" (or "full direct product") of the family, sometimes written $\prod_{\alpha\in A} E_\alpha$.

2b) But there is also a special vector subspace V of $\prod_{\alpha\in A} E_\alpha$. It consists of those elements $x := (x_\alpha)_{\alpha\in A}$, where of course $x_\alpha \in E_\alpha$ for each α, for which $x_\alpha = 0_\alpha$, the zero vector in E_α, except for *finitely many* indices α:

$$V := \left\{x \in \prod\nolimits_{\alpha\in A} E_\alpha : \#\{\alpha \in A : x_\alpha \neq 0_\alpha\} < \infty\right\}.$$

Notice that the set $\{\alpha \in A : x_\alpha \neq 0_\alpha\}$ will usually be different for different $x \in V$. V is called the "restricted direct product" or "direct sum" of the family, and sometimes denoted $\bigoplus_{\alpha\in A} E_\alpha$. (The word "product" and the sign $\prod$ in the one case, and "sum" and $\bigoplus$ in the other, are due to category-theoretical considerations we need not go into here. Indeed, in categorical terms, $\bigoplus$ is a coproduct.)

3) If F is a vector subspace of E, define an equivalence relation $\sim$ on E by agreeing that $x \sim y$ means $x - y \in F$. (It is trivial that this is an equivalence relation.) The set of equivalence classes becomes a vector space if we set

$$\lambda[x] + \mu[y] := [\lambda x + \mu y],$$

where $[a]$ denotes the equivalence class of $a \in E$. The resulting vector space is the *quotient space of* E *by* F, usually denoted E/F. The mapping $E \longrightarrow E/F : x \mapsto [x]$ is called the *projection* of E on to F.

Definition 8.1. Let E_1, E_2 be vector spaces over $\mathbb{K}$. A mapping $T : E_1 \longrightarrow E_2$ is called a *linear map* or *linear transformation* if

$$(\forall \lambda, \mu \in \mathbb{K})(\forall x, y \in E_1) \quad T(\lambda x + \mu y) = \lambda T(x) + \mu T(y).$$

It is a curiosity of mathematical notation that $T(x)$ is often abbreviated to Tx in this context. [I suspect this is because of the close relation with matrices and their multiplication.]

If $E_1 = E_2 = E$, T is often called a *linear operator in* E. If $E_2 = \mathbb{K}$, T is called a *linear functional on* E. (This curious name was invented before the theory was clarified. A "functional" was a scalar-valued mapping whose domain was a set of *functions*, in the days when the idea of a "function" was still restricted to things that had as domain a subset of $\mathbb{R}^n$ or $\mathbb{C}^n$. The name has survived because it is a convenient way of making the distinction.)

$\mathcal{L}(E_1; E_2)$ is the set of all linear maps from E_1 to E_2. In category theory it might be denoted $\mathrm{Hom}(E_1, E_2)$, "Hom" being the set of morphisms in the category of vector spaces.

Lemma 8.2. *If* E_1, E_2 *are vector spaces over* $\mathbb{K}$*, then* $\mathcal{L}(E_1; E_2)$ *is also a vector space over* $\mathbb{K}$*, with the "pointwise linear operations"*

$$(\forall T_1, T_2 \in \mathcal{L}(E_1; E_2))(\forall \lambda, \mu \in \mathbb{K}) \quad (\lambda T_1 + \mu T_2)(x) := \lambda T_1(x) + \mu T_2(x). \qquad \square$$

It is worth pointing out that the *commutativity* of $\mathbb{K}$ is essential here.

Definition 8.3. In particular, the vector space $\mathcal{L}(E; \mathbb{K})$ is called the *algebraic dual* (or algebraic conjugate) of E. We shall denote it, for the moment, by E^*.

The dual itself has a dual $(E^*)^*$, which we denote E^{**} and call the *algebraic bidual* or *algebraic second dual* or *algebraic second conjugate space* of E.

There is a natural mapping $J : E \longrightarrow E^{**}$, given by

$$(\forall \phi \in E^*)(\forall x \in E) \quad (J(x))(\phi) := \phi(x). \tag{25}$$

It is easy to check that $J(x)$ is indeed a linear functional on E^*, i.e. an element of E^{**}.

For some purposes it is helpful to write $\phi(x)$ as $\langle x, \phi \rangle$; then $\langle , \rangle$ is a mapping $E \times E^* \longrightarrow \mathbb{K}$ referred to as the *dual pairing*. There is a second dual pairing $\langle , \rangle^* : E^* \times E^{**} \longrightarrow \mathbb{K}$, and then the definition (25) is equivalent to

$$\langle \phi, J(x) \rangle^* = \langle x, \phi \rangle. \tag{26}$$

Lemma 8.4. *The bidual map* $J : E \longrightarrow E^{**}$ *is linear and one-one. It is onto if and only if* E *is finite-dimensional.*

Proof. The linearity of J follows from the definition of the linear operations in the dual spaces. To prove it is one-one, we only need to show that $J(x) = 0 \Longrightarrow x = 0$.

So, suppose $x \in E$ and $J(x) = 0$. This means that, for any $\phi \in E^*$, $\langle \phi, J(x) \rangle^* = 0$, or, by (26), $\langle x, \phi \rangle = 0$. Thus, we can conclude that $x = 0$ if we can show that

$$\text{if } 0 \neq x \in E \text{, there exists } \phi \in E^* \text{ such that } \phi(x) \neq 0. \tag{27}$$

(This is argument by the contrapositive.)

A. Suppose, by 7.1, that B is a basis in E. If $0 \neq x \in E$, then x is a linear combination of elements of B; suppose, in this linear combination, $b \in B$ has non-zero coefficient λ. Define ϕ_x by specifying its values on the elements of B: $\phi_x(b) = 1$, and $\phi_x(c) = 0$ if $b \neq c \in B$. The linear map $\phi_x : E \longrightarrow \mathbb{K}$ is constructed by linear extension from its values on B, and $\phi_x(x) = \lambda \neq 0$. This establishes (27), and so J is one-one. [This could be proved slightly more simply by constructing B to include $\{x\}$. But see below.]

Let $\{e_1, e_2, \ldots, e_n\}$ be a basis in E, if it is finite-dimensional. There is a so-called *dual basis* $\{e_1^*, e_2^*, \ldots, e_n^*\}$ in E^*, where the functional e_r^* is specified by

$$e_r^*(e_s) = \delta_{rs} \quad \text{(the Kronecker } \delta\text{: } \delta_{rs} = 0 \text{ when } s \neq r,\ \delta_{rr} = 1\text{)}. \tag{28}$$

If $\phi \in E^*$, then $\phi = \phi(e_1)e_1^* + \phi(e_2)e_2^* + \cdots + \phi(e_n)e_n^*$, because they agree on every element of the basis in E; so the e_r^* span E^*, and are easily seen to linearly independent too. Thus they form a basis in E^*, whose dimension is therefore also n. But then E^{**} must also have dimension n. Hence, as J is one-one and $\dim E = \dim E^{**}$, J must be onto.

B. Now suppose that E is *not* finite-dimensional. Recall the basis B for E used above, and consider the set N of those linear functionals ϕ on E with the property that $\{b \in B : \phi(b) \neq 0\}$ is finite. N is a vector subspace of E^*, but it is not all of E^*, since there is a linear functional on E that takes the value 1 on all elements of B. Now N has a basis B_0^*, by 7.1, which is included in a basis B^* for the whole of E^*, again by by 7.1; and we may define $\Phi \in E^{**}$ by setting it to be 0 on $B_0^* = N \cap B^*$ and to be 1 on $B^* \setminus N$, and then extending it linearly to all of E^*.

Suppose $\Phi = J(x)$ for some $x \in E$. Then, by (26), $\langle x, \phi \rangle = \langle \phi, \Phi \rangle^*$ for all $\phi \in E^*$. In particular, for any $\phi \in N$, $\langle x, \phi \rangle = \langle \phi, \Phi \rangle^* = 0$ by construction. However, construct ϕ_x as at A above. $\phi_x \in N$, since it was constructed to be non-zero on exactly one element "*b*" of B, and $\langle x, \phi_x \rangle \neq 0$. The contradiction shows that Φ cannot be in the image of J. $\square$

In fact, if the dimension of E is an *infinite* cardinal $\aleph$, the dimension of E^* is $2^{\aleph}$. This is not very difficult to prove, but requires some knowledge of the theory of cardinal numbers.

Linear functionals are of central importance, though it is difficult at first sight to see why. Functional analysis began as an attempt to generalize well-known facts about linear transformations of finite-dimensional vector spaces to infinite dimensions; and this entailed the formulation of linear algebra in a basis-free fashion. The dual space is the "abstract" way of describing the space of "row vectors" if the original space consists of "column vectors". When ξ is a $1 \times n$ row vector and x is an $n \times 1$ column vector, the matrix product ξx is 1×1, and may be regarded as a scalar; thus we get the dual pairing that I denoted $\langle x, \xi \rangle$. The basis in the space of row vectors that is dual to the standard basis in the space of column vectors (see (28)) is just the standard basis in the space of row vectors. But this easy correspondence does not work if you start from a non-standard basis for the column vectors.

However, if E is infinite-dimensional, we have just seen that the algebraic dual of E is usually absurdly large; for this and other reasons, infinite-dimensional vector spaces are often specified with a topology, and one considers the *continuous* (or *topological*) *dual* instead.

Definition 8.5. A *topological vector space* (over $\mathbb{K}$), also called a *linear topological space*, is a vector space E over $\mathbb{K}$, together with a topology on E such that both the linear operations

$$E \times E \longrightarrow E : (x, y) \mapsto x + y\,, \quad \mathbb{K} \times E \longrightarrow E : (\lambda, x) \mapsto \lambda x$$

are continuous (the products are given the product topology).

A *metric linear space* is a topological linear space whose topology is defined by a specific metric. A topological vector space is *metrizable* if its topology *may* be defined by a metric. It is *normable* if its topology may be defined by a metric derived from a norm; and it is a *normed space* if it has a specified norm that defines its topology. It is easily seen that a normed space is a topological vector space with respect to the topology defined by the norm.

In practice, the distinction between a "metrizable" (or "normable") space and a "metric" (or "normed") one is often treated rather casually.

Examples 8.6. The examples that follow are not intended to be treated very rigorously.

(i) $C(\Omega; \mathbb{K})$, where Ω is a compact topological space, is a normed space. See 4.3. The linear operations are defined "pointwise", and the norm is the "supremum norm", whose characteristic property is that $f_n \to f$ in norm means "$f_n \to f$ uniformly on Ω".

(ii) l^p, the space of $\mathbb{K}$-valued sequences $(\xi_1, \xi_2, \xi_3, \dots)$ such that $\sum |\xi_n|^p < \infty$. Here the index $p > 0$. The linear operations are "coordinatewise".

If $1 \le p < \infty$, l^p is normed by $\|(\xi_1, \xi_2, \xi_3, \dots)\| := (\sum |\xi_n|^p)^{1/p}$.

This expression is obviously "definite" and "homogeneous". That it is "subadditive" is not obvious (it is a version of Minkowski's inequality). The case we shall be strongly interested in is $p = 2$, for which subadditivity is easy and well-known.

If $0 < p < 1$, give l^p a metric: if $x := (\xi_1, \xi_2, \xi_3, \dots)$ and $y := (\eta_1, \eta_2, \eta_3, \dots)$, both in l^p, set

$$d(x, y) := \sum |\xi_n - \eta_n|^p\,,$$

which must also converge and is indeed a metric making l^p a topological vector space.****

(iii) l^∞, which consists of the *bounded* $\mathbb{K}$-valued sequences with supremum norm.

(iv) Let c_{00} consist of *finitely non-zero* $\mathbb{K}$-valued sequences, $(\xi_1, \xi_2, \dots, \xi_n, 0, 0, \dots)$ (the index n after which all terms are zero is allowed to vary with the sequence under consideration), with coordinatewise operations. It is possible to define a topology on c_{00} with the characteristic property that $x_n := (\xi_1^{(n)}, \xi_2^{(n)}, \xi_3^{(n)}, \dots)$ converges to $x := (\xi_1, \xi_2, \xi_3, \dots)$ if and only if there is some $k \in \mathbb{N}$ such that $\xi_r = 0$ and $\xi_r^{(n)} = 0$ for all n and for all $r > k$, and $\xi_r^{(n)} \to \xi_r$ as $n \to \infty$ for each $r \le k$.

(v) Let s be the vector space of *all* $\mathbb{K}$-valued sequences, with the termwise operations. It may be topologized (as a topological vector space) by the product topology on $\prod_{n=1}^{\infty} \mathbb{K}$; I have so far avoided discussing this topology, but in this case it may be defined by a metric:

$$d((\xi_1, \xi_2, \xi_3, \dots), (\eta_1, \eta_2, \eta_3, \dots)) := \sum_{n=1}^{\infty} \frac{1}{2^n} \frac{|\xi_n - \eta_n|}{1 + |\xi_n - \eta_n|}\,.$$

(The same topology may also be defined by very many other metrics; for instance, instead of the numerical coefficients 2^{-n}, one might have any sequence (α_n) of positive numbers such that $\sum \alpha_n < \infty$. The function $t/(1+t)$ may also be substituted by other functions. One metric often given is $d'((\xi_1, \xi_2, \xi_3, \dots), (\eta_1, \eta_2, \eta_3, \dots)) := \sum_{n=1}^{\infty} 2^{-n} \min(1, |\xi_n - \eta_n|)$.)

The characteristic property of the topology defined by this metric is that a sequence (x_n) in s, where $x_n := (\xi_1^{(n)}, \xi_2^{(n)}, \xi_3^{(n)}, \dots)$, converges to $x := (\xi_1, \xi_2, \xi_3, \dots)$ if and only if $\xi_k^{(n)} \to \xi_k$ as $n \to \infty$ for each separate k. ****

(vi) If $f : \mathbb{R} \longrightarrow \mathbb{R}$, we say that f is *of compact support* if there is a compact set K in $\mathbb{R}$ such that $f(x) = 0$ when $x \notin K$. Let $\mathcal{D}(\mathbb{R})$ be the space of C$^\infty$ functions $\mathbb{R} \longrightarrow \mathbb{R}$ which are of compact support, with pointwise linear operations. There is a topology on $\mathcal{D}(\mathbb{R})$ with the characteristic property that $f_n \to f$ in $\mathcal{D}(\mathbb{R})$ if and only if there are a compact set K in $\mathbb{R}$ and a natural number k such that $f(x) = 0$ and $f_n(x) = 0$ for any n when $x \notin K$, and $f_n^{(r)} \to f^{(r)}$ as $n \to \infty$ uniformly on K for every $r \geq 0$. (Here $f^{(r)}$ is the rth derivative.)

This is reminiscent of *(iv)*. In both cases the topology is in fact the *inductive limit* of topologies on subspaces. Its formal definition is rather complicated, and I shan't give it; but, in particular, it is *not* metrizable. However, this is an important example, because it is the beginning of Laurent Schwartz's theory of "distributions" [not to be confused with probability distributions!], which were invented to make sense of things like the Dirac δ-function and have had very serious applications to partial differential equations.

(vii) Let G be an open set in $\mathbb{C}$. Let $A(G)$ be the vector space of functions $G \longrightarrow \mathbb{C}$ that are holomorphic on G, with pointwise linear operations. A set $U \subseteq A(G)$ is open if, for any $f \in U$, there are some compact subset K of G and $\epsilon > 0$ such that

$$\{g \in A(G) : (\forall x \in K) \;\; |f(x) - g(x)| < \epsilon\} \subseteq U .$$

This topology is metrizable, and its characteristic property is that $f_n \to f$ if and only if, for any compact $K \subseteq G$, $f_n|K \to f|K$ uniformly. We may meet it later.

Whilst the above examples are all interesting in some way, the sequence spaces *(ii)*–*(v)* may seem a little artificial. But, just as *(iv)* has an analogous function space *(vi)*, and *(iii)* is vaguely related to *(i)*, there are other function spaces that are of great practical importance. I shall not try to give a completely rigorous treatment of them, but I think it is worthwhile to describe them in approximate terms, as they are the reason for much of the theory.

Example 8.7. Let U be a geometrically simple region in $\mathbb{R}^n$, let $p \geq 1$, and let $L^p_{\mathbb{K}}(U)$ be the class of functions $f : U \longrightarrow \mathbb{K}$ which are "integrable in pth power" (in modulus, that is): $\displaystyle\int_U |f(x)|^p \, dx < \infty$, where the integral is an n-dimensional volume integral. Again, this is a complex vector space under pointwise operations; a norm on it may be defined by

$$\|f\|_p := \left(\int_U |f(x)|^p \, dx \right)^{1/p} . \tag{29}$$

Then $L^p_{\mathbb{K}}(U)$ is a Banach space, often called the *Lebesgue space with exponent* p.

Unfortunately, what I have just written is not entirely true. If we restrict attention to bounded uniformly continuous functions and to a bounded region U, the integrals definitely exist, but the space is *not* complete. It is possible to give it a "completion" by a rather simple abstract construction (one adds points corresponding to Cauchy sequences that do not converge), but, although one obtains a Banach space by this means, it is, obviously, no longer a space of functions on U, but a space of equivalence classes of sequences of functions.

The problem arises because the "limit", in any appropriate sense, of a Cauchy sequence (f_n) (with respect to the norm (29)) may not itself be integrable in the Riemann sense. This was, in retrospect, one of the reasons why the Riemann integral had to be generalized. The Lebesgue integral, on the other hand, enables one to integrate any function that is at all likely to have an integral; and if one employs the Lebesgue integral, any sequence of functions that are Lebesgue-integrable in pth power and form a $\|\|_p$-Cauchy sequence $\|\|_p$-converges to a function in $L^p_{\mathbb{K}}(U)$.

However, the enlargement of the class of functions that may be integrated introduces another lesser difficulty — some non-zero nonnegative functions have zero integrals, so that the formula (29) does not define a *norm* (it is no longer definite) but only a seminorm. To overcome this objection, one considers the elements of $L^p(U)$ as *equivalence classes* of functions, where the equivalence relation is "equality almost everywhere", or, in effect,

$$f \sim g \quad \text{means} \quad \int_U |f(x) - g(x)|^p \, dx = 0\,.$$

After all this, one does obtain a genuine Banach space. Its elements are not, strictly speaking, functions, but equivalence classes of functions. The theorem that $L^p_{\mathbb{K}}(U)$, understood in this sense, is complete is often called the Riesz-Fischer theorem. (Riesz and Fischer proved it independently of each other around 1907. I read somewhere that they met for the first time more than twenty-five years later.)

The elements of $L^p_{\mathbb{K}}(U)$ are commonly denoted as, and thought of as, functions, albeit with occasional acknowledgement that two of them are "equal" when the functions by which they are represented are equal almost everywhere.

In several of the spaces I have listed, the real version also has an order structure, such as we exploited in proving the Stone-Weierstrass theorem. In particular, it is sometimes a *vector lattice* (crudely speaking, the "maximum" of two elements of the space is also in the space). But notice that there are topological vector spaces that are not metrizable (that is, their topologies cannot be defined by a metric), and a metrizable t.v.s. need not be normable. A simple example is the space s of *(v)* above. It is easily proved that any open set which contains the origin in s includes a whole infinite-dimensional linear subspace. In a normed space, the "open unit ball" $\{x \in E : \|x\| < 1\}$ is an open set containing the origin, but it cannot include even a one-dimensional linear subspace.

Definition 8.8. Let E be a vector space over $\mathbb{K}$. A metric d on E is *translation-invariant* (or just *invariant*) if $d(x+z, y+z) = d(x,y)$ for all $x, y, z \in E$.

A metric derived from a norm is necessarily invariant, and the other metrics described above also are. An example of a non-invariant metric on $\mathbb{R}$ is, for instance,

$$d(x,y) := \left| \frac{x}{1+|x|} - \frac{y}{1+|y|} \right| ;$$

it defines the same topology as the usual metric.

Theorem 8.9. *A topological vector space E is metrizable if and only there is a* sequence (U_n) *of open sets in E such that*

(a) $\{0\} = \bigcap_{n=1}^{\infty} U_n$, *and*

(b) *for any open set* U *such that* $0 \in U$ *, there is some* n *for which* $U_n \subseteq U$ *.*
If the topology of E *is metrizable, then it may be defined by an invariant metric.* □

The second half of this theorem (or rather a generalization of it to topological groups) is sometimes called the *Birkhoff-Kakutani theorem*. I omit the proof. But, as a consequence, we shall always assume any metric that we use on a metrizable linear space is invariant.

Lemma 8.10. *If two invariant metrics on a vector space* E *define the same topology, and* E *is complete with respect to one of them, it is complete with respect to the other.* □

(In fact, both metrics will define the same Cauchy sequences. This is certainly not true for non-invariant metrics.)

Definition 8.11. Let E be a vector space over $\mathbb{K}$. An *inner product* on E is a function

$$E \times E \longrightarrow \mathbb{K} : (x, y) \mapsto \langle x, y \rangle$$

such that the following conditions are satisfied for all $x, y, z \in E$ and $\lambda, \mu \in \mathbb{K}$:

$$\begin{array}{llll} \textit{(a)} & \langle x, y \rangle = \overline{\langle y, x \rangle}\,, & \textit{(b)} & \langle \lambda x + \mu y, z \rangle = \lambda \langle x, z \rangle + \mu \langle y, z \rangle\,, \\ \textit{(c)} & \langle x, x \rangle \geq 0\,, & \textit{(d)} & \langle x, x \rangle = 0 \Longrightarrow x = 0\,. \end{array}$$

(If $\mathbb{K} = \mathbb{R}$, the conjugation in *(a)* is to be ignored. When $\mathbb{K} = \mathbb{C}$, *(a)* ensures that $\langle x, x \rangle$ is real, so that *(c)* makes sense; and *(a)* and *(b)* imply that $\langle x, \lambda y + \mu z \rangle = \overline{\lambda} \langle x, y \rangle + \overline{\mu} \langle x, z \rangle$. Thus $\langle , \rangle$ is linear in the first argument and conjugate-linear in the second, or *sesquilinear*.)

The pair $(E, \langle , \rangle)$ is called an *inner product space*.

If $p = 2$ in 8.7, $L^2_{\mathbb{K}}(U)$ is an inner product space, with

$$\langle \{f\}, \{g\} \rangle := \int_\Omega f(x) \overline{g(x)}\, \mu(dx)$$

as the inner product (I write $\{f\}$ for the element of $L^2_{\mathbb{K}}(U)$ which is the equivalence class of the function f, but the distinction is, as I said, mostly not observed). The other standard examples such as $\mathbb{K}^n$ and $l^2_{\mathbb{K}}$ are, in effect if not in strict logic, special cases of this.

Lemma 8.12. *Let* E *be an inner product space. Then the formula* $\|x\| := \sqrt{\langle x, x \rangle}$ *defines a norm on* E*. Furthermore, one has the Cauchy-Schwarz inequality:* $|\langle x, y \rangle| \leq \|x\| \|y\|$ *.* □

An inner product space is always understood to be furnished with this norm, and in this sense *an inner product space is a normed space.* But not every norm arises from an inner product. Less obviously, there are normable spaces for which *no* norm describing the given topology can be derived from an inner product.

Hence, there are topological vector spaces, *some* of which are metrizable (have topologies that may be defined by a metric); amongst metrizable topological vector spaces, *some* are normable; and amongst normable topological vector spaces, *some* can be given an inner product whose associated norm defines the correct topology.

§9. Completeness properties.

The notion of completeness may be extended to topological vector spaces in general, but I shall only discuss it for metric linear spaces with invariant metrics. Indeed, the results I shall give do not generalize very easily to larger classes of space, and I think the versions I shall give are the "natural" ones, however much they may have been extended since.

Definition 9.1. A complete normed space, as already mentioned, is called a *Banach* space.

An inner product space that is complete with respect to the norm defined by the inner product is a *Hilbert space*.

A complete metric linear space (with invariant metric) is a *Fréchet* space.

[A Fréchet space is often defined as a complete metric *locally convex* space. We shall have little to say about local convexity, which is not needed in the theorems below.]

Definition 9.2. Let E be a vector space. A set $A \subseteq E$ is *balanced* if $\lambda a \in A$ whenever $a \in A$ and $|\lambda| \leq 1$, where $\lambda \in \mathbb{K}$. A is *absorbent* (in E) if, for any $x \in E$, there exists some $\delta > 0$ such that $\lambda x \in A$ whenever $|\lambda| < \delta$.

"Balanced" is the same as other authors would call "star-shaped and circled".

It is often convenient to use some abbreviated notations. If A and B are subsets of a vector space E, we write $-A$ for $\{-a : a \in A\}$, $A + B$ for $\{a + b : a \in A \ \& \ b \in B\}$ and, similarly, $A - B$ for $\{a - b : a \in A \ \& \ b \in B\}$; and, if $\lambda \in \mathbb{K}$, we write λA for $\{\lambda a : a \in A\}$. It is important to realize that, as these are operations on *subsets* of E, the "rules of algebra" do not necessarily apply. For instance, $A + A$ and $2A$ are in principle not the same (although $2A \subseteq A + A$) and $A - A$ is $\{0\}$ only if A is a singleton. But it is true that $A + (-A) = A - A$. Notice too that, if $A = \emptyset$, both $A + B$ and λA are $\emptyset$.

If A is balanced, then, in particular, $A = -A$. If E is a topological vector space, A is open, and $\lambda \neq 0$, then λA is open. ****

Lemma 9.3. *Let E be a topological vector space.*

(a) *Any open subset of E that contains the origin is absorbent.*

(b) *If U is an open subset of E that contains the origin, there is a balanced open subset V that also contains the origin, and is such that $V + V \subseteq U$.*

Proof. Let U be open, $0 \in U$. Take any $x \in E$; $0x = 0 \in U$, the "zero" on the *right* being the zero *vector*. Scalar multiplication is continuous at $(0, x) \in \mathbb{K} \times E$, so there exist $\delta > 0$ and an open set $V_1 \ni x$ such that

$$|\lambda| < \delta \ \& \ y \in V_1 \Longrightarrow \lambda y \in U,$$

and, in particular, $|\lambda| < \delta \Longrightarrow \lambda x \in U$. So U is absorbent.

As addition is continuous at $(0, 0) \in E \times E$, there is an open set $W_1 \ni 0$ such that

$$x, y \in W_1 \Longrightarrow x + y \in U;$$

as scalar multiplication is continuous at $(0, 0) \in \mathbb{K} \times E$, there exist an open set $W \ni 0$ and a number $\epsilon > 0$ such that, for $\lambda \in \mathbb{K}$ and $x \in E$,

$$|\lambda| < \epsilon \ \& \ x \in W \Longrightarrow \lambda x \in W_1.$$

Define $V := \bigcup_{\mu \in \mathbb{K},\ |\mu| \leq \epsilon} \mu W$. Then V is open and balanced, and $V \subseteq W_1$, and

$$V + V \subseteq W_1 + W_1 \subseteq U. \qquad \square$$

Proposition 9.4. *Let* $T : E \longrightarrow F$ *be a linear mapping of topological vector spaces. Suppose that* $T(E)$ *is of the second Baire category in* F*. Then, for any open set* U *of* E *with* $0_E \in U$*,* $0_F \in \mathrm{int}_F(\mathrm{cl}_F(T(U)))$.

Proof. By 9.3, there is a balanced open set $V \ni 0_E$ in E such that $V + V \subseteq U$ (and so $V - V \subseteq U$), and V is absorbent. Take any $\beta > 1$ in $\mathbb{R}$. As V is absorbent, there exists, for any $x \in E$, some $n \in \mathbb{N}$ such that $\beta^{-n}x \in V$, or $x \in \beta^n V$. That is, $E = \bigcup_{n=1}^{\infty} \beta^n V$. Therefore, T being linear,

$$T(E) = T\Big(\bigcup\nolimits_{n=1}^{\infty} \beta^n V\Big) = \bigcup\nolimits_{n=1}^{\infty} \beta^n T(V).$$

But $T(E)$ is not of the first category, so at least *one* of the balanced sets $\beta^n T(V)$ must have a closure with non-empty interior in F. The same must hold for $T(V)$, so there are a point $x_0 \in \mathrm{cl}_F(T(V))$ and an open set $W \ni x_0$ such that $W \subseteq \mathrm{cl}_F(T(V))$, and then

$$\begin{aligned} 0_F \in W - W \subseteq \mathrm{cl}_F(T(V)) - \mathrm{cl}_F(T(V)) &\subseteq \mathrm{cl}_F(T(V) - T(V)) \\ &\text{(by continuity of subtraction in } F\text{)****} \\ = \mathrm{cl}_F(T(V - V)) \subseteq \mathrm{cl}_F(T(U)). & \end{aligned}$$

But $W - W = \bigcup_{x \in W} (W - x)$ is an open set. □

In fact, the topology of E is not really essential here; we just need a balanced absorbent set V such that $V + V \subseteq U$. The substance of the argument is in F.

Theorem 9.5. (The open mapping theorem.) *Suppose that* E *is a Fréchet space in the sense of* 9.1*, and* F *is a linear metric space. If* $T : E \longrightarrow F$ *is a continuous linear map such that* $T(E)$ *is of the second category in* F*, then* T *is open.*

Proof. For each $m \in \mathbb{N}$, there is by 9.4 some $\rho'_m > 0$ such that

$$\mathrm{cl}_F(T(B_E(0; 2^{-m}))) \supseteq B_F(0; \rho'_m).$$

Let $\rho_1 := \rho'_1$, and if $\rho_1, \rho_2, \dots, \rho_m$ have been defined, let $\rho_{m+1} := \min(\rho'_{m+1}, \frac{1}{2}\rho_m)$. Then $\rho_1 > \rho_2 > \cdots > \rho_m \downarrow 0$ as $m \to \infty$, and $B_F(0; \rho_m) \subseteq \mathrm{cl}_F(T(B_E(0; 2^{-m})))$ for each m.

Take any $y \in B_F(0; \rho_m)$. Write $y_m := y$. Thus $y_m \in \mathrm{cl}_F(T(B_E(0; 2^{-m})))$. By 3.3,

$$B_F(y_m; \rho_{m+1}) \cap T(B_E(0; 2^{-m})) \neq \emptyset,$$

and there exists $x_m \in B_E(0; 2^{-m})$ such that $Tx_m \in B_F(y_m; \rho_{m+1})$. Hence

$$y_{m+1} := y_m - Tx_m \in B_F(0; \rho_{m+1}).$$

If we have already chosen $y_{m+k} \in B_F(0; \rho_{m+k}) \subseteq \mathrm{cl}_F(T(B_E(0; 2^{-m-k})))$, take in the same way $x_{m+k} \in B_E(0; 2^{-m-k})$ such that $Tx_{m+k} \in B_F(y_{m+k}; \rho_{m+k+1})$, and let $y_{m+k+1} := y_{m+k} - Tx_{m+k}$. In this way we define sequences $x_m, x_{m+1}, x_{m+2}, \dots$ and $y_m, y_{m+1}, y_{m+2}, \dots$ by induction, so that, for each $n \geq m$, $x_n \in B_E(0; 2^{-n})$ and $y_{n+1} = y_n - Tx_n \in B_F(0; \rho_{n+1})$.

For each $n \geq m$, let $\sigma_n := x_m + x_{m+1} + \cdots + x_n$. Then, whenever $p \geq n \geq m$,

$$\begin{aligned} d_E(\sigma_n, \sigma_{n+1}) = d(0, \sigma_{n+1} - \sigma_n) = d(0, x_{n+1}) < 2^{-n-1}, \\ d_E(\sigma_n, \sigma_p) < 2^{-n-1} + 2^{-n-2} + \cdots + 2^{-p} < 2^{-n}. \end{aligned} \tag{30}$$

Thus $(\sigma_n)_{n\geq m}$ is a Cauchy sequence in E, and must converge in E, let us say to σ; and $d_E(\sigma,\sigma_m) = \lim_{n\to\infty} d_E(\sigma_n,\sigma_m) \leq 2^{-m}$, so that — recall $\sigma_m = x_m \in B_E(0;2^{-m})$ —

$$d_E(\sigma,0) \leq d_E(x_m,0) + d_E(\sigma,x_m) < 2^{-m+1}, \quad \sigma \in B_E(0;2^{-m+1})\,. \tag{31}$$

On the other hand,

$$y - T(\sigma_n) = y_m - Tx_m - Tx_{m+1} - \cdots - Tx_n = y_n - Tx_n \in B_F(0;\rho_{n+1})\,,$$

and so $T(\sigma_n) \to y$ as $n \to \infty$. As T is continuous and $\sigma_n \to \sigma$, $T(\sigma_n) \to T\sigma$ too. [This is the first time we have appealed to the continuity of T.] So $T\sigma = y$.

Recall that y was any point of $B_F(0;\rho_m)$, and that $\sigma \in B_E(0;2^{-m+1})$ by (31). So we have shown that $B_F(0;\rho_m) \subseteq T(B_E(0;2^{-m+1}))$. Hence, for any $z \in E$,

$$B_F(Tz;\rho_m) = Tz + B_F(0;\rho_m) \subseteq Tz + T(B_E(0;2^{-m+1})) = T(B_E(z;2^{-m+1}))\,, \tag{32}$$

for any $m \in \mathbb{N}$. This is the main task completed; now we finish off.

Let U be any open set in E, and suppose that $z \in U$. There exists $m(z) \in \mathbb{N}$ (depending on z) such that $B_E(z;2^{-m(z)+1}) \subseteq U$, by 3.6; and, by (32),

$$T(U) \supseteq T(B_E(z;2^{-m(z)+1})) \supseteq B_F(Tz;\rho_{m(z)})\,. \tag{33}$$

Hence $T(U) = \bigcup_{z\in U} B_F(Tz;\rho_{m(z)})$, which is an open set in F. □

Corollary 9.6. *With the above hypotheses, necessarily* $T(E) = F$.

Proof. Indeed, $T(E)$ is an open set containing 0, so is absorbent in F by 9.3. So, for any $y \in F$, there is some $\lambda \in \mathbb{K} \setminus \{0\}$ such that $\lambda y \in T(E)$, and $\lambda y = Tx$ for some $x \in E$. Thus $y = T(\lambda^{-1}x) \in T(E)$. □

Thus, if E is Fréchet and $T(E)$ is of the second category in the linear metric space F, a continuous linear mapping T must in fact be open and onto F. This forces the topology of F to be the *quotient topology* defined by T (I shan't go into the details here), and therefore:

Corollary 9.7. *With the same hypotheses, F must be complete.*

Proof. Suppose that (y_n) is a Cauchy sequence in F. By passing to a subsequence (if necessary), we may assume that $d_F(y_n,y_{n+1}) < \rho_n$ for each n, $y_{n+1} \in B_F(y_n;\rho_n)$.

As T is surjective by 9.6, there exists $x_1 \in E$ such that $Tx_1 = y_1$. If x_n has been chosen so that $Tx_n = y_n$, choose $x_{n+1} \in B_E(x_n;2^{-n+1})$ so that $Tx_{n+1} = y_{n+1}$, using (32). Thus, inductively, one forms a sequence (x_n) such that $d_E(x_n,x_{n+1}) < 2^{-m+1}$ and $Tx_n = y_n$ for each n. As at (30), (x_n) is Cauchy, so it converges to $x \in E$. As T is continuous, $y_n = Tx_n \to Tx$. This suffices (if a subsequence of a Cauchy sequence converges, the whole sequence converges to the same limit). □

For convenience, I have referred to (32), but the result could be formulated in the form: *if $T : E \longrightarrow F$ is a continuous and open map between linear metric spaces, and E is complete, then so is F.* It is only necessary to choose each ρ_m so that (32) is satisfied.

Although this result is not without interest, one wishes commonly to prove that T is open by means of 9.5; and, to show $T(E)$ is of the second category, one would normally use 5.4

and prior knowledge that F is complete. [There are metric linear spaces that are of the second category without being complete.] In fact, the most often quoted version of 9.5 is

Note 9.8. A continuous linear surjection of Fréchet spaces is open.

Corollary 9.9. *Suppose that $T : E \longrightarrow F$ is a continuous linear mapping between linear metric spaces, where E is complete and $T(E)$ of second category in F; and suppose T is one-one. Then $T(E) = F$, and the inverse mapping $T^{-1} : F \longrightarrow E$ which is, therefore, defined, is also continuous.* □

This result is interesting from the viewpoint of category theory. A linear mapping between vector spaces which is one-one and onto has an inverse *mapping* (of sets), which is also linear. A homomorphism of groups which is one-one and onto has an inverse mapping (of sets) which is also a homomorphism of groups. But a continuous mapping between topological spaces which is one-one and onto need not have a continuous inverse. 9.9 shows that this peculiarity cannot occur for continuous linear mappings between Fréchet spaces.

Definition 9.10. Suppose that $f : A \longrightarrow B$ is any mapping of sets. The *graph* of f is the subset $G(f) := \{(x, f(x)) : x \in A\}$ of $A \times B$.

According to the nowadays generally accepted definition of a mapping, the graph of f is, therefore, just f itself. It is, nevertheless, fairly standard to speak in this context of the graph of f to emphasize that one is thinking of a subset of $A \times B$.

If E and F are metric linear spaces, then $E \times F$ is a vector space (see the reamrks at the beginning of §5) and has a metric given by (13). It is easy to see that $E \times F$ becomes a metric linear space, and that it is complete if E and F are both complete.

Theorem 9.11. (The closed graph theorem.) *Let E and F be Fréchet spaces in the sense of* 9.1, *and let $T : E \longrightarrow F$ be a linear map. Then T is continuous if and only if its graph is a closed subset of the metric linear space $E \times F$.*

Proof. Assume T continuous. If $(x, y) \notin G(T)$, then $Tx \neq y$. Let $\epsilon := \frac{1}{2} d_F(y, Tx) > 0$. By continuity of T, there exists $\delta > 0$ such that $T(B_E(x; \delta)) \subseteq B_F(Tx; \epsilon)$. Let $\kappa := \min(\delta, \epsilon)$, and suppose that $(a, b) \in B_{E\times F}((x, y); \kappa)$; then $d_E(a, x) + d_F(b, y) < \kappa$, so that $d_F(b, y) < \epsilon$ and $d_E(a, x) < \delta$, and, consequently, $d_F(Ta, Tx) < \epsilon$. However,

$$2\epsilon = d_F(y, Tx) \leq d_F(Ta, Tx) + d_F(b, Ta) + d_F(y, b) < 2\epsilon + d_F(b, Ta),$$

from which $Ta \neq b$, or $(a, b) \notin G(T)$. Hence, any point not in $G(T)$ belongs to an open set (here $B_{E\times F}((x, y); \kappa)$) which does not meet $G(T)$. This proves that $G(T)$ is closed. [This result holds for any continuous map from a topological space to a Hausdorff topological space; it does not require linearity or a metric. Indeed, the general proof is easier.****]

Now suppose that $G(T)$ is closed in $E \times F$. It is clear that $G(T)$ is a vector subspace of $E \times F$, and, as $E \times F$ is complete, so is $G(T)$ as a metric linear space in its own right. The projection on the first coordinate, $\pi_1 : G(T) \longrightarrow E : (x, Tx) \mapsto x$, is trivially continuous, one-one, and onto E. By 9.9, its inverse $\gamma : x \mapsto (x, Tx) : E \longrightarrow G(T)$ is also continuous. The projection $\pi_2 : G(T) \longrightarrow F : (x, Tx) \mapsto Tx$ is also continuous. Thus the composition $\pi_2 \circ \gamma : x \mapsto Tx : E \longrightarrow F$ is continuous, as required. □

The closed graph theorem is important because it suggests that some results might be obtained on the hypothesis that the graph is closed, even though the mapping is not continuous or defined on the whole of E. Although we shall not discuss problems of this kind, they are very important, for instance in quantum mechanics, where it is essential to allow differentiation operators.

The definition 1.8 of equicontinuity may be given a slightly different form; it is not necessary that Ω or Ψ be a metric space.

Definition 9.12. Let Ω be a topological space, F a topological vector space, and suppose that $\{f_\alpha : \alpha \in A\}$ is a family of mappings $f_a : \Omega \longrightarrow E$. The family is described as *equicontinuous at* $x \in \Omega$ if, for any open set U of F containing 0, there is an open set V of Ω containing x such that, for every $\alpha \in A$, $f_\alpha(V) \subseteq f_\alpha(x) + U$ (or, equivalently, $V \subseteq f_\alpha^{-1}(f_\alpha(x) + U)$).

This agrees with 1.8 if Ω is metric and E is a metric linear space with invariant metric.

Proposition 9.13. *Let E, F be topological vector spaces, and suppose that E is of second category in itself. For $n \in \mathbb{N}$, let $f_n : E \longrightarrow F$ be a continuous linear map. Suppose that, for each $x \in E$, the sequence $f_n(x)$ in F converges as $n \to \infty$ to $f(x) \in F$, where $f : E \longrightarrow F$ is linear. Then, for every $x \in E$, the family (f_n) is equicontinuous at x.*

As a consequence, f is also continuous.

Proof. Firstly, it will suffice to prove that, for any open set U of F containing 0, there is an open set V of E containing 0 such that $f_n(V) \subseteq U$ (or $V \subseteq f_n^{-1}(U)$) for each n. For then

$$f_n(x + V) = f_n(x) + f(V) \subseteq f_n(x) + U \quad \text{for each } n \text{ and each } x \in E.$$

So take such a U. By 9.3, there is an open set U_1 of E such that $0 \in U_1$ and $U_1 + U_1 \subseteq U$, and a balanced open set U_2 such that $U_2 + U_2 \subseteq U_1$.

Suppose that $x \notin U_1$. If $a \in (x + U_2) \cap U_2$, $a = x + u$ for some $u \in U_2$, and $x = a - u \in U_2 + U_2 \subseteq U_1$ (as U_2 is balanced). As this is not true, necessarily $(x + U_2) \cap U_2 = \emptyset$, and so $x \notin \mathrm{cl}_F(U_2)$. Thus $F \setminus U_1 \subseteq F \setminus \mathrm{cl}_F(U_2)$, or $\mathrm{cl}_F(U_2) \subseteq U_1$.

Take $y \in E$. Then $f(y) \in F$, and, again by 9.3, there exists $\epsilon > 0$ such that

$$f(\lambda y) = \lambda f(y) \in U_2 \quad \text{when } |\lambda| < \epsilon.$$

For such a λ, $f_n(\lambda y) \to f(\lambda y)$, and, as U_2 is open and contains $f(\lambda y)$, there is some N such that $n \geq N + 1 \Longrightarrow f_n(\lambda y) \in U_2$.

On the other hand, $f_n^{-1}(U_2)$ is an open set in E (and contains 0) for each $n \in \mathbb{N}$, so that $\bigcap_{n=1}^{N} f_n^{-1}(U_2)$ is also an open set in E containing 0 and must be absorbent: there exists $\delta > 0$ such that $\lambda y \in \bigcap_{n=1}^{N} f_n^{-1}(U_2)$ if $|\lambda| < \delta$.

Putting the arguments of the preceding two paragraphs together,

$$|\lambda| < \min(\delta, \epsilon) \Longrightarrow \lambda y \in \bigcap_{n=1}^{\infty} f_n^{-1}(U_2) \subseteq M := \bigcap_{n=1}^{\infty} f_n^{-1}(\mathrm{cl}_F(U_2)).$$

Thus, M is absorbent in E, and $\bigcup_{n=1}^{\infty} nM = E$. It is also closed (as each $f_n^{-1}(\mathrm{cl}_F(U_2))$ is).

Now, E is of second category; so there must exist some $n \in \mathbb{N}$ such that $\mathrm{cl}_E(nM)$ has nonempty interior. But $\mathrm{cl}_E(nM) = n\,\mathrm{cl}_E(M)$, and so $M = \mathrm{cl}_E(M)$ itself has non-empty interior W. I assert that $W - W$ is an open set containing the origin in E. Indeed, for any $y \in W$, $W - y$ is open and contains 0, and $V := W - W = \bigcup_{y \in W} (W - y)$ is itself open. However, for any $n \in \mathbb{N}$, the definitions give

$$\begin{aligned} f_n(V) &= f_n(W) - f_n(W) \subseteq f_n(M) - f_n(M) \\ &\subseteq \mathrm{cl}_F(U_2) - \mathrm{cl}_F(U_2) \subseteq U_1 - U_1 \subseteq U. \end{aligned}$$

This therefore concludes the proof that (f_n) is equicontinuous.

It remains to show that f must be continuous. Suppose again that U is any open set in F containing 0, and let U_1 be a balanced open set such that $U_1 + U_1 \subseteq U$. By equicontinuity, there is an open set W in E, containing 0, such that $f_n(W) \subseteq U_1$ for all $n \in \mathbb{N}$. Suppose $y = f(x)$ for some $x \in W$; then $f_n(x) \to f(x)$, and so there exists $N \in \mathbb{N}$ such that $f_N(x) - f(x) \in U_1$ and $y = f(x) \in f_N(x) - U_1 \subseteq U_1 + U_1 \subseteq U$. Consequently, $f(W) \subseteq U$. This proves that f is continuous (at the origin, and so everywhere). □

The above result is one of several similar theorems — similar, that is, in their proofs, but differing in their hypotheses. The general idea is perhaps due to Steinhaus, although it was Saks who noticed the applicability of Baire's theorem.

§10. Normed spaces.

There are many results which assume a really satisfactory form only for normed spaces.

Definition 10.1. Let E and F be normed spaces, and $T : E \longrightarrow F$ a linear map. T is said to be *bounded* if there exists $K \geq 0$ such that $\|Tx\|_F \leq K\|x\|_E$ for all $x \in E$. Such a constant K is called a *bound* for T.

It is usually not necessary to distinguish in the notation between the norms in different spaces, as I just did; for the context ought to remove all ambiguity.

Lemma 10.2. *If E and F are normed, the* bounded *linear transformations $E \longrightarrow F$ form a vector subspace of $\mathcal{L}(E;F)$.*

Proof. Suppose S and T are both bounded with respect to the norms in E and F, with bounds α and β. Then

$$\begin{aligned} (\forall x \in E) \quad \|(\lambda S + \mu T)x\| &= \|\lambda(Sx) + \mu(Tx)\| \quad \text{by definition} \\ &\leq |\lambda|\|Sx\| + |\mu|\|Tx\| \\ &\leq (|\lambda|\,\alpha + |\mu|\,\beta)\|x\|. \end{aligned} \tag{34}$$

So $|\lambda|\,\alpha + |\mu|\,\beta$ is a bound for $\lambda S + \mu T$. □

Definition 10.3. Given normed spaces E and F, the vector space of all bounded linear transformations from E to F will be denoted $L(E;F)$.

Lemma 10.4. *A linear transformation $T : E \longrightarrow F$ is bounded if and only if it is continuous.* □

It is customary to denote by 0 not only the scalar 0 and the vector 0, but also the 'trivial' vector space with only one element (namely the zero vector), and the zero linear mapping which carries every vector of its domain into the zero vector of the codomain. The *linear* *mapping* 0 is the *vector* 0 in $\mathcal{L}(E;F)$ or in $L(E;F)$. Whilst these conventions are obviously undesirable in strict logic, the context always removes any ambiguity.

I have *not* shown either that there are circumstances in which $L(E;F) \neq 0$, i.e. that there may exist a nonzero bounded linear map from E to F, or that there may be an unbounded linear map, $L(E;F) \neq \mathcal{L}(E;F)$. We shall see that, provided E and F are not trivial, neither is $L(E;F)$, and perhaps also that $L(E;F) = \mathcal{L}(E;F)$ if and only if *either* F is trivial *or* E is finite-dimensional.

Definition 10.5. If $T : E \longrightarrow F$ is a bounded linear map between normed spaces E and F, the infimum of the bounds for T is called the *uniform operator-norm* of T (with respect to the given norms in E and F), or just the norm of T, and may be denoted $\|T\|_{E;F}$ or $\|T\|_{L(E;F)}$ or simply $\|T\|$.

Lemma 10.6. *If* $E \neq 0$ *and* $T \in \mathcal{L}(E;F)$, *then* T *is bounded if and only if* $\{\|Tx\|_F : \|x\|_E = 1\}$ *is bounded; the operator-norm of* T *is itself a bound for* T*; and indeed, if* $E \neq 0$,

$$\begin{aligned}\|T\|_{E;F} &:= \inf\{K : (\forall x \in E)\ \ \|Tx\|_F \le K\|x\|_E\} = \sup\left\{\frac{\|Tx\|_F}{\|x\|_E} : 0 \neq x \in E\right\} \\ &= \sup\{\|Tx\|_F : \|x\|_E = 1\} = \sup\{\|Tx\|_F : \|x\|_E \le 1\}. \qquad \square\end{aligned}$$

If $E = 0$, the only possible linear operator out of E is the zero operator; its norm is 0. The second and third formulæ above are wrong in that case.

Lemma 10.7. *In the above circumstances,* $\|\|_{E;F}$ *is a norm in* $L(E;F)$.

Proof. In the calculation (34) above, one may take $\alpha := \|S\|$ and $\beta := \|T\|$, $\lambda = \mu = 1$. It follows that $\|S\| + \|T\|$ is a bound for $S + T$, and, therefore, that $\|S + T\| \le \|S\| + \|T\|$. It is trivial that $\|\lambda S\| = |\lambda|\|S\|$ and that $\|S\| = 0$ if and only if S is the zero map. $\square$

An important consequence of 10.6 is that $\|Tx\|_F \le \|T\|_{E,F}\|x\|_E$ for all $x \in E$ and $T \in L(E;F)$.

Definition 10.8. Suppose that E, F, G are vector spaces. A mapping $T : E \times F \longrightarrow G$ is called *bilinear* if, for every $a \in E$, the mapping $F \longrightarrow G : y \mapsto T(a, y)$ is linear, and, for every $b \in F$, the mapping $E \longrightarrow G : x \mapsto T(x, b)$ is linear.

There are similar definitions of trilinear, quadrilinear, $\ldots$, n-linear mappings (or maps, or transformations; they are not usually called "operators").

Definition 10.9. Suppose that E, F, G are normed spaces. A non-negative number β is a *bound* for the bilinear mapping $T : E \times F \longrightarrow G$ when $\|T(x, y)\|_G \le \beta\|x\|_E\|y\|_F$ for all $x \in E$ and $y \in F$.

There is a similar definition for multilinear mappings of higher "degrees". It is easily checked that the set $\mathcal{L}(E, F; G)$ of bilinear maps $E \times F \longrightarrow G$ (more generally, the set $\mathcal{L}(E_1, E_2, \ldots, E_n;G)$ of n-linear maps $E_1 \times E_2 \times \cdots \times E_n \longrightarrow G$) is a vector space over $\mathbb{K}$ with respect to pointwise operations, and that the bounded bilinear (or n-linear) mappings form a vector subspace $L(E, F; G)$ (or $L(E_1, E_2, \ldots, E_n; G)$). If any of the spaces E_i or G

is trivial, then $L(E_1, E_2, \ldots, E_n; G)$ is also trivial, but otherwise the formula

$$\|T\| := \sup\left\{ \frac{\|T(x_1, x_2, \ldots, x_n)\|_G}{\|x_1\|_{E_1}\|x_2\|_{E_2}\cdots\|x_n\|_{E_n}} : 0 \neq x_1 \in E_1, \ldots, 0 \neq x_n \in E_n \right\}$$

defines a norm in $L(E_1, E_2, \ldots, E_n; G)$ that may be called the multilinear operator norm.

Lemma 10.10. $T \in \mathcal{L}(E_1, E_2, \ldots, E_n; G)$ *is continuous (as a mapping of the product space* $E_1 \times E_2 \times \cdots \times E_n$ *, with product topology, into* F*) if and only if it is bounded.* □

Spaces of multilinear operators are, vaguely speaking, superfluous. Let me explain why.

Definition 10.11. Let Ω, Ψ be metric spaces. A mapping $f : \Omega \longrightarrow \Psi$ is *isometric* or an *isometry* if, for any $x, y \in \Omega$, $d_\Psi(f(x), f(y)) = d_\Omega(x, y)$. Similarly, if E and F are normed spaces, a mapping $f : E \longrightarrow F$ is isometric if, for all $x, y \in E$,

$$\|f(x) - f(y)\|_F = \|x - y\|_E .$$

There is a simple theorem of Mazur that any isometric mapping between normed spaces over $\mathbb{R}$ must be linear, but it is too long to be given here***. At any rate, an *isometric isomorphism* between normed spaces over $\mathbb{K}$ is a linear isomorphism which is also an isometry. Clearly its inverse will also be a linear isometry.

In the following Lemma, "natural" means, informally, "independent of all internal choices that might be used in its construction but are not indicated in the statement". Although this is the basic meaning, it is rather vague. The category theorists, wishing to avoid all mention of the internal structure of the objects they deal with — which is where choices might commonly arise —, have introduced a definition of *natural transformations of functors*. They would say that $\mathcal{L}(\,\cdot\,;\,\cdot\,)$ is a *functor* of two arguments, contravariant in the first argument and covariant in the second [note that differential geometers use "contravariant" and "covariant" the wrong way round!], and that the "natural isomorphisms" below are also natural transformations of functors in this sense. (The category-theoretical definition does not capture the intuitive idea very well, because there may even be natural transformations of functors in the category-theoretical sense that *do* involve choices, but it is mathematically precise.) For simplicity, I shall treat only the case $n = 2$, from which the general case can be obtained by a straightforward induction.

Lemma 10.12. *For any vector spaces* E, F, G*, there is a natural isomorphism*

$$\Phi : \mathcal{L}(E; \mathcal{L}(F; G)) \longrightarrow \mathcal{L}(E, F; G) .$$

If E, F, G *are normed, then* $L(E; L(F; G)) \subseteq \mathcal{L}(E; \mathcal{L}(F; G))$ *and*

$$\Phi(L(E, L(F, G))) = L(E, F; G) ,$$

and $\Phi|L(E; L(F; G) : L(E; L(F; G) \longrightarrow L(E, F; G)$ *is a natural isomorphism that is also an isometry with respect to the operator-norms.*

Proof. Define $\Phi : \mathcal{L}(E; \mathcal{L}(F; G)) \longrightarrow \mathcal{L}(E, F; G)$ by setting, given $T \in \mathcal{L}(E; \mathcal{L}(F; G))$,

$$(\forall x \in E)(\forall y \in F) \quad \Phi(T)(x, y) = T(x)(y) . \tag{35}$$

(Explanation: $T(x) \in \mathcal{L}(F; G)$ by definition, so may act on y.) It is easily checked that

$\Phi(T)$ is bilinear, so that Φ does map $\mathcal{L}(E;\mathcal{L}(F;G))$ into $\mathcal{L}(E,F;G)$. Moreover, (35) may be reversed: for any $S \in \mathcal{L}(E,F;G)$ and any $x \in E$, define $\Psi(S)(x)$ by

$$(\forall y \in F) \quad \Psi(S)(x)(y) = S(x,y), \tag{36}$$

and then $\Psi(S)x \in \mathcal{L}(F;G)$ for each x and $\Psi(S)$ is in $\mathcal{L}(E;\mathcal{L}(F;G))$. Hence, Ψ and Φ are mutually inverse one-to-one correspondences, and they are trivially linear, so they are isomorphisms. They are obviously natural.

If either E or F is trivial, all four spaces under consideration are trivial. Otherwise, if T is bounded with respect to norms in E, F, and G, (35) gives

$$\|\Phi(T)(x,y)\|_G = \|T(x)(y)\|_G \le \|T(x)\|_{L(F,G)}\|y\|_F \le \|T\|_{L(E;L(F;G))}\|x\|_E\|y\|_F\,, \tag{37}$$

so that $\Phi(T) \in L(E,F;G)$, and $\|\Phi(T)\| \le \|T\|$. In a similar way, (36) gives

$$\|\Psi(S)(x)(y)\|_G \le \|S\|_{L(E,F;G)}\|x\|_E\|y\|_F\,,$$

which tells us firstly that $\Psi(S)(x)$ is bounded for each x, with $\|\Psi(S)(x)\|_{L(F;G)} \le \|S\|\|x\|$, and then that $\|\Psi(S)\| \le \|S\|$. Writing $T := \Psi(S)$, we deduce that $\|T\| \le \|\Phi(T)\|$, and, with (37) shown previously, we derive the result. $\square$

Definition 10.13. If E, F are vector spaces over $\mathbb{K}$ and $T \in \mathcal{L}(E,E,\ldots,E;F)$, where E is repeated n times, then T is described as *symmetric* (a symmetric n-linear mapping) if, for every permutation σ of $\{1,2,\ldots,n\}$, and for every choice of $x_1,x_2,\ldots,x_n \in E$,

$$T(x_1,x_2,\ldots,x_n) = T(x_{\sigma(1)},x_{\sigma(2)},\ldots,x_{\sigma(n)})\,.$$

In the same way, T is *skew-symmetric* if, for every choice of σ and of $x_1,x_2,\ldots,x_n \in E$,

$$T(x_1,x_2,\ldots,x_n) = \epsilon(\sigma)\,T(x_{\sigma(1)},x_{\sigma(2)},\ldots,x_{\sigma(n)})\,,$$

where $\epsilon(\sigma)$ denotes the sign of the permutation σ, as usual.

For simple examples, suppose $E = \mathbb{K}^2$, $F = \mathbb{K}$, and $n = 2$. Then the mapping

$$\mathbb{K}^2 \times \mathbb{K}^2 \longrightarrow \mathbb{K} : ((\xi_1,\xi_2),(\eta_1,\eta_2)) \mapsto \xi_1\eta_2 + \xi_2\eta_1$$

is symmetric bilinear, and the mapping

$$((\xi_1,\xi_2),(\eta_1,\eta_2)) \mapsto \xi_1\eta_2 - \xi_2\eta_1$$

is skew-symmetric bilinear. (Notice that I have *not* written $\mathbb{K}^2 \times \mathbb{K}^2$ as $\mathbb{K}^4$, to avoid any suspicion that we are treating it as the vector space $\mathbb{K}^2 \oplus \mathbb{K}^2$). It is clear that what we are talking about in this very elementary case are symmetric polynomials and skew-symmetric polynomials of degree 1 in each variable (that is the meaning of bilinearity) in two variables (because the dimension of $\mathbb{K}^2$ is 2).

Lemma 10.14. *Let E, F, G be normed spaces. Then the mapping*

$$L(E;F) \times L(F;G) \longrightarrow L(E;G)$$

given by composition, $(S,T) \mapsto T \circ S$, is bounded bilinear, and its bilinear operator norm *(with respect to the operator-norms in $L(E,F)$ and $L(F,G)$) does not exceed* 1*. That is,*

$$\|T\circ S\|_{L(E;G)} \leq \|T\|_{L(F;G)}\|S\|_{L(E;F)}\,. \qquad \square$$

In fact, if E,F,G are non-trivial, the bilinear operator norm of the composition is exactly 1, but we cannot yet prove it.

Lemma 10.15. *Let E be a normed space and F a Banach space. Then $L(E;F)$ is a Banach space with respect to the operator-norm.*

Proof. Let (T_n) be a Cauchy sequence in operator-norm. Given $\epsilon > 0$, there exists $N(\epsilon)$ such that $n,m \geq N(\epsilon) \Longrightarrow \|T_n - T_m\| < \epsilon$.

Take some fixed vector $x \in E$. If $x = 0$, then

$$T_n x = 0 \to 0 \quad \text{as} \quad n \to \infty\,.$$

If $x \neq 0$,

$$n,m \geq N(\epsilon/\|x\|) \Longrightarrow \|T_n x - T_m x\| = \|(T_n - T_m)x\| \leq \|T_n - T_m\|\|x\| < \epsilon\,, \tag{38}$$

so that $(T_n x)$ is a Cauchy sequence in F. As such, it again converges in F. Let us define Tx to be the limit in F of the sequence $(T_n x)$, for *any* choice of x in E. T is our *candidate* for the limit of the sequence (T_n).

Firstly, $T : E \longrightarrow F$ is a linear transformation. Indeed,

$$\begin{aligned}\lambda Tx + \mu Ty &= \lambda \lim T_n x + \mu \lim T_n y = \lim(\lambda T_n x + \mu T_n y)\\ &= \lim T_n(\lambda x + \mu y) = T(\lambda x + \mu y)\,.\end{aligned}$$

Secondly, T is bounded. For this, use 10.6. If $y \in E$ and $\|y\| = 1$, then by (38)

$$n,m \geq N(\epsilon) \Longrightarrow \|T_n y - T_m y\| < \epsilon\,.$$

In particular, $n,m \geq N(1) \Longrightarrow \|T_n y\| < \|T_m y\| + 1$. Now fix $m = N(1)$. There is a constant $K := \left\|T_{N(1)}y\right\| + 1$ such that $\|T_n y\| < K$ for $n \geq N(1)$; hence,

$$\|Ty\| = \|\lim(T_n y)\| = \lim \|T_n y\| \leq K$$

and, as this is true for any $y \in E$ such that $\|y\| = 1$, $T \in L(E;F)$ and $\|T\| \leq K$ by 10.6. Our "candidate" T indeed is an object of the right kind (belongs to the right party).

It remains to check that $T_n \to T$ in the desired sense. Given $\epsilon > 0$ and $y \in E$ such that $\|y\| = 1$, (38) shows $\|T_n y - T_m y\| < \epsilon$ when $n,m \geq N(\epsilon)$. Hence, if $n \geq N(\epsilon)$,

$$\|T_n y - Ty\| = \lim_{m\to\infty} \|T_n y - T_m y\| \leq \epsilon\,,$$

and 10.6 shows that $\|T_n - T\| \leq \epsilon$. This shows that $T_n \to T$ in operator-norm. $\square$

Notice that the proofs that $T \in L(E;F)$ and that $T_n \to T$ are very similar.

Lemma 10.16. *Let E be a normed space, and let F be a closed vector subspace of F. Then the quotient space E/F (see* 3 *of* §8*) is also a normed space, with the "quotient norm"*

$$\|[x]\|_{E/F} := \inf\{\|y\|_E : y \in E \;\;\&\;\; y - x \in F\} = \inf\{\|y\|_E : y \in [x]\} \tag{39}$$

for any $x \in E$. Furthermore, if E is complete with respect to $\|.\|_E$, then so is E/F with *respect to* $\|.\|_{E/F}$ *; and, if* $F \neq E$ *, the projection* $\pi : E \longrightarrow E/F : x \mapsto [x]$ *is a bounded linear map of norm* 1.

Proof. That $\|.\|_{E/F}$ is a seminorm is easily checked, and of course from (39)

$$\|\pi(x)\|_{E/F} = \|[x]\|_{E/F} \leq \|x\|_E \tag{40}$$

for any $x \in E$. If $\|[x]\|_{E/F} = 0$, then there exist elements $z_n \in E$ such that $x - z_n \in F$ and $\|z_n\| < 1/n$ (say). Hence, $x - z_n \to x$, and, as F is closed, $x \in F$, or $[x] = 0_{E/F}$. So $\|.\|_{E/F}$ is a norm in E/F.

Now, suppose that $F \neq E$, and let $[x] \neq 0_{E/F}$. As just shown, $\|[x]\|_{E/F} > 0$. Let $\epsilon > 0$. From (39), there exists $x' \in E$ such that $[x'] = [x]$ and $\|x'\|_E < (1+\epsilon)\|[x]\|_{E/F}$. So $\|\pi(x')\|_{E/F} = \|[x]\|_{E/F} > (1+\epsilon)^{-1}\|x'\|_E$. This shows $\|\pi\| > (1+\epsilon)^{-1}$, and (40) shows that $\|\pi\| \leq 1$; hence, $\|\pi\| = 1$ exactly.

Let $([x_n])$ be a Cauchy sequence in E/F. To show that it converges, it will be sufficient to show it has a convergent subsequence. By passing to a subsequence if necessary, and then renumbering, we may assume that

$$\|[x_n - x_{n+1}]\|_{E/F} = \|[x_n] - [x_{n+1}]\|_{E/F} < 2^{-n} \quad \text{for each } n. \tag{41}$$

Define $x'_1 := x_1$. If x'_n has been defined so that $[x'_n] = [x_n]$, then definition (39), with (41), shows that there exists $y_{n+1} \in E$ such that

$$y_{n+1} - (x_{n+1} - x'_n) \in F \text{ and } \|y_{n+1}\|_E < 2^{-n}.$$

Then define $x'_{n+1} := y_{n+1} + x'_n$, and it follows that $[x'_{n+1}] = [x_{n+1}]$ and

$$\|x'_{n+1} - x'_n\| < 2^{-n} \quad \text{for each } n.$$

The sequence (x'_n) is Cauchy in E, for, if $m \geq n$,

$$\|x'_m - x'_n\|_E \leq \|x'_m - x'_{m-1}\|_E + \cdots + \|x'_{n+1} - x'_n\|_E < 2^{-m+1} + \cdots + 2^{-n} < 2^{-n+1}.$$

Hence there exists some $x \in E$ such that $\|x'_n - x\|_E \to 0$ as $n \to \infty$, and

$$\|[x_n] - [x]\|_{E/F} \leq \|x'_n - x\|_E \to 0.$$

This proves convergence of $([x_n])$ in E/F. $\square$

The completeness proof is essentially the same as at 9.7.

Remark 10.17. It is immediately apparent that any vector subspace F of a normed space E is also a normed space (one defines $\|x\|_F := \|x\|_E$ for any $x \in F$). If E is a Banach space, then F is a Banach space if and only if it is closed in E. Similarly, if E_1, E_2 are normed spaces with norms $\|.\|_1, \|.\|_2$, the direct sum $E_1 \oplus E_2$ (see 2) of §5) may be normed by

$$\|(x,y)\|_\oplus := \|x\|_1 + \|y\|_2.$$

Then $E_1 \oplus E_2$ is a Banach space with respect to $\|\|_\oplus$ if and only if both E_1 and E_2 are Banach with respect to their norms. Compare (13) of 3.12.

Definition 10.18. Let E be a topological vector space. Its *topological dual*, or *continuous dual*, or *dual* or *conjugate space*, is the set of continuous linear functionals $E \longrightarrow \mathbb{K}$. I shall denote it by E'. Thus, if E is a normed space, E' consists of all the *bounded* linear functionals on E. [Recall from 10.4 that a linear functional is bounded if and only if it is continuous.]

Recall that E^* denoted the set of *all* linear functionals $E \longrightarrow \mathbb{K}$, 8.1; and so $E' \subseteq E^*$. Granted the existence of an algebraic basis in E, we showed $E^* \neq 0$. However, there are topological vector spaces whose topological duals are trivial. Here is an example.

Example 10.19. Suppose $0 < p < 1$, and let $\mathcal{L}^p(0,1)$ denote the space of functions $(0,1) \longrightarrow \mathbb{R}$ such that the improper Riemann integral

$$\int_0^1 |f(t)|^p \, dt := \lim_{\substack{b\uparrow 1\\ a\downarrow 0}} \int_a^b |f(t)|^p \, dt$$

is defined (i.e. is finite). Then $\mathcal{L}^p(0,1)$ is a pseudometric linear space if we set ******

$$d(f,g) := \int_0^1 |f(t) - g(t)|^p \, dt \,.$$

[Notice we do not take the "*p*th root" here; this pseudometric does not come from a norm.]

Suppose, if possible, that ϕ is a continuous nonzero linear functional on $\mathcal{L}^p(0,1)$. Then there is some $f \in \mathcal{L}^p(0,1)$ such that $\phi(f) \neq 0$. Multiplying by a suitable constant, we may suppose that $\phi(f) \geq 1$.

For $0 \leq \tau \leq 1$, define f^τ to be the function $(0,1) \longrightarrow \mathbb{R}$ defined by

$$f^\tau(t) := \begin{cases} f(t)\,, & \text{if } t \leq \tau\,, \\ 0\,, & \text{if } t > \tau\,. \end{cases}$$

Thus f^0 is the zero-function and $f^1 = f$, and $\tau \mapsto f^\tau$ is a continuous mapping $[0,1] \longrightarrow \mathcal{L}^p(0,1)$. So $\phi(f^\tau)$ is a continuous real-valued function, whose value at 0 is 0 and at 1 is 1. Now

$$d(0, f^\tau) = \int_0^1 |f^\tau(t)|^p \, dt = \int_0^\tau |f(t)|^p \, dt$$

is a continuous function of τ that increases steadily from 0 to $d(0,f)$, so there must be some value of τ for which $d(0,f^\tau) = \frac{1}{2} d(0,f) = d(f^\tau, f) = d(0, f - f^\tau)$. However,

$$1 = \phi(f) = \phi(f^\tau) + \phi(f - f^\tau)\,,$$

so that either $\phi(f^\tau) \geq \frac{1}{2}$ or $\phi(f - f^\tau) \geq \frac{1}{2}$ (or both). Take $f_1 := 2f^\tau$ if $\phi(f^\tau) \geq \frac{1}{2}$, and otherwise let $f_1 := 2(f - f^\tau)$. In this way, given f such that $\phi(f) = 1$, we obtain f_1 such that $\phi(f_1) \geq 1$ and $d(0,f_1) = 2^{p-1} d(0,f)$. Let $f_2 := (f_1)_1$ and so on inductively; we get f_n such that $\phi(f_n) \geq 1$ and $d(0,f_n) = 2^{n(p-1)} d(0,f)$. Hence, $f_n \to 1$ according to the metric d, and yet $\phi(f_n) \not\to 0$. This is impossible if ϕ is to be continuous, and the only flaw in the argument is the initial assumption that $\phi(f) \neq \emptyset$.

The conclusion must be that the only continuous linear functional on the pseudometric linear space $\mathcal{L}^p(0,1)$ is zero.

We shall see in due course that this behaviour is impossible in normed spaces; in fact, it is impossible in *locally convex* spaces, which is why local convexity is an important condition.

Definition 10.20. Let E be a topological vector space, and $a \in E$. A *basic weakly open neighbourhood* of a in E is a set of the form

$$N(a; \phi_1, \phi_2, \ldots, \phi_n; \epsilon_1, \epsilon_2, \ldots, \epsilon_n) := \{x \in E : |\phi_i(x-a)| < \epsilon_i \quad \text{for} \quad 1 \leq i \leq n\}$$

for some $n \in \mathbb{N}$, $\epsilon_1, \epsilon_2, \ldots, \epsilon_n > 0$, and $\phi_1, \phi_2, \ldots, \phi_n \in E'$. A set $U \subseteq E$ is *weakly open* if, for every $a \in U$, there is a basic weakly open neighbourhood of x that is included in U.

The *weak topology* on E, often denoted $\sigma(E, E')$, is the class of weakly open sets in E.

Remark 10.21. That the weak topology on E is a topology with respect to which E is a topological vector space is an *exercise*. We have not developed the theory of t.v.s. sufficiently to make it as trivial as it should be.

The notation $\sigma(E, E')$ is appropriate, because we could define a similar weak topology $\sigma(E, G)$ on E for any vector subspace G of E^*. However, the "weak topology" on a given t.v.s E is always understood to be $\sigma(E, E')$ unless there is some indication otherwise. By definition, the linear functionals in E' are continuous with respect to the given topology on E; $\sigma(E, E')$ is the weakest (that is, the smallest) topology on E that still makes E a topological vector space and these linear functionals continuous.

A sequence (x_n) in E converges weakly to $x \in E$, i.e. converges with respect to the weak topology, if and only if $\phi(x_n) \to \phi(x)$ for every $\phi \in E'$. (The notation $x_n \rightharpoonup x$ is sometimes found to express weak convergence.) Part of the importance of dual spaces lies in the possibility of using weak convergence to construct solutions of various problems.

Definition 10.22. Given E as above, the *weak* topology* on E' (also denoted $\sigma(E', E)$) is the topology in which a set U is open (and said to be *weak* open*) if, for any $\alpha \in U$, there is a set of the form

$$N^*(\alpha; x_1, x_2, \ldots, x_n; \epsilon_1, \epsilon_2, \ldots, \epsilon_n) := \{\phi \in E' : |(\phi - \alpha)(x_i)| < \epsilon_i \quad \text{for} \quad 1 \leq i \leq n\}$$

for some $n \in \mathbb{N}$, $\epsilon_1, \epsilon_2, \ldots, \epsilon_n > 0$, and $x_1, x_2, \ldots, x_n \in E$, that is included in U.

This fits the remark in 10.21, since E can be regarded as a subspace of $(E')^*$.

There are several theorems that, broadly speaking, state ways in which weak topologies resemble finite-dimensional topologies. A very notable example is the following.

Theorem 10.23. (Alaoglu-Bourbaki.) *Let E be a normed space. Then the closed unit ball in the dual, $\{\phi \in E' : \|\phi\| \leq 1\}$, is $\sigma(E', E)$-compact (or, as it is often put, weak*-compact).*

Proof. For any $r \in \mathbb{R}$, let $D(r) := \{x \in \mathbb{K} : |x| \leq r\}$. This is a compact subset of $\mathbb{K}$. Now consider the class Ω of *all* functions (not necessarily linear or even continuous) $\psi : E \longrightarrow \mathbb{K}$ such that, for each $x \in E$, $|\psi(x)| \leq \|x\|$. This class is $\prod_{x \in E} D(\|x\|)$ (see the discussion of the Multiplicative Axiom in §7A). However, Tychonov's theorem (again, see §7B) says that $\Omega = \prod_{x \in E} D(\|x\|)$ is compact in the product topology. (See the following remark for the long-delayed discussion of this product topology.)

Let $\alpha, \beta \in \mathbb{K}$ and $x, y \in E$. The maps $\Omega \longrightarrow \mathbb{K}$ given by $\psi \mapsto \psi(x)$, $\psi \mapsto \psi(y)$, and $\psi \mapsto \psi(\alpha x + \beta y)$ are continuous. So must be

$$f_{\alpha,\beta,x,y} : \psi \mapsto \alpha\psi(x) + \beta\psi(y) - \psi(\alpha x + \beta y)\,.$$

Hence $f^{-1}_{\alpha,\beta,x,y}(\{0\})$ is a closed subset of Ω, as must be the intersection of all such subsets,

$$\Omega_0 := \bigcap\nolimits_{\alpha,\beta\in\mathbb{K},\, x,y\in E} f^{-1}_{\alpha,\beta,x,y}(\{0\})\,.$$

By 3.16*(a)*, Ω_0 is also compact as a subset of Ω with the product topology.

On the other hand, Ω_0 is, by definition, precisely the set of elements of Ω that are linear maps $E \longrightarrow \mathbb{K}$. And, by the definition also, $\psi \in \Omega$ only if $|\psi(x)| \le \|x\|$ for all $x \in E$; thus Ω_0 is, in fact, the closed unit ball in E'.

Finally, the weak* topology on the closed unit ball in E' is the restriction to Ω_0 of the product topology in Ω (see below). □

Remark 10.24. The proof above is the one given by Bourbaki. Alaoglu's theorem was restricted to Hilbert space and its proof was superficially quite different; and there are other proofs which are in a sense generalizations of Alaoglu's.

I have avoided discussing the product topology on product spaces with infinitely many factors, but something must be said here. Suppose that, for each β in an index set B, Ω_β is a topological space. A *base* for the topology $\mathcal{T}_\beta$ of Ω_β is a subclass $\mathcal{B}_\beta$ of $\mathcal{T}_\beta$ such that any member of $\mathcal{T}_\beta$ is a union of members of $\mathcal{B}_\beta$.

Let $x := (x_\beta)_{\beta\in B}$ be an element of $\Omega := \prod_{\beta\in B}\Omega_\beta$ (in effect, x is a function defined on B whose value at β is $x_\beta \in \Omega_\beta$, the "βth coordinate" of x; see §7A). Given a finite subset F of B and, for each $\beta \in F$, $U_\beta \in \mathcal{B}_\beta$ such that $x_\beta \in U_\beta$, the corresponding *basic product neighbourhood* of x in Ω is the set of all elements $(y_\beta)_{\beta\in B}$ of Ω such that $y_\beta \in U_\beta$ for each $\beta \in F$. This basic product neighbourhood of x may be thought of as the product $\prod_{\beta\in B} V_\beta$, where $V_\beta := U_\beta$ for $\beta \in F$ and $V_\beta := \Omega_\beta$ when $\beta \notin F$. A set A is defined to be open in the product topology for Ω if, for every $x \in A$, there is some basic product neighbourhood Q for x in Ω such that $Q \subseteq A$. The class of open sets defined by this prescription does not depend on the choice of the bases $\mathcal{B}_\beta$ for the topologies $\mathcal{T}_\beta$, and is indeed a topology in Ω.

The definition of the product topology ensures that the *projections* $\pi_\beta : \Omega \longrightarrow \Omega_\beta$ are all continuous, where the βth projection π_β carries each point $x := (x_\beta)$ of Ω into its βth coordinate x_β. Indeed, if U is open in Ω_β, its inverse image is itself a basic product neighbourhood of each of its points (take $F := \{\beta\}$ and $U_\beta := U$). The product topology is the smallest topology that makes all the projections continuous.

Now suppose that Ω_0 is a subset of a topological space Ω. The *subspace topology* on Ω_0 consists of all the subsets of Ω_0 that are of the form $U \cap \Omega_0$, where U is open in Ω. An alternative formulation is that the open sets of Ω_0 are the inverse images $i^{-1}(U)$ of open sets U of Ω under the *inclusion map* $i : \Omega_0 \longrightarrow \Omega : x \mapsto x$. It is trivial that this does define a topology in Ω_0, which is the weakest topology such that i is continuous.

The assertion in the proof of 10.23 that the weak* topology on the closed unit ball in E' is exactly the subspace topology on Ω_0 induced from the product topology on Ω results from chasing the definitions above.

Since $B' := \{\phi \in E' : \|\phi\| \le 1\}$ is weak*-compact and the mappings $\phi \mapsto \phi(x)$ (for $x \in E$) are continuous $E' \longrightarrow \mathbb{K}$ when E' has the weak* topology, we have a mapping of E into $C(\Omega_0;\mathbb{K})$ for a certain compact topological space Ω_0, namely $\Omega_0 := B'$. This mapping is linear, but *so far we do not even know whether* Ω_0 *has more than one point*, namely the zero linear functional.

§11. The Hahn-Banach theorem

Definition 11.1. Let E be a vector space over $\mathbb{R}$. A *sublinear functional* on E is a mapping $p : E \longrightarrow \mathbb{R}$ such that

$$\text{(a)} \quad (\forall x, y \in E) \quad p(x+y) \le p(x) + p(y) \quad \text{and}$$
$$\text{(b)} \quad (\forall x \in E)(\forall \lambda \in \mathbb{R}) \quad \lambda \ge 0 \Longrightarrow p(\lambda x) = \lambda p(x)\,.$$

Condition *(a)* states that p is *subadditive*, and condition *(b)* that it is *positive homogeneous*. (In fact, *(b)* will follow from *(a)* and the weaker condition that $\lambda > 0 \Longrightarrow p(\lambda x) = \lambda p(x)$.***) A norm on E is a sublinear functional, but a sublinear functional need not be a norm.

Definition 11.2. Let E be a vector space over $\mathbb{K}$. The *codimension* in E of a subspace F of E is the algebraic dimension of E/F. F is *finite-codimensional* [or *of finite codimension*] if its codimension is finite, and is a *hyperplane* if its codimension is 1. [It should be admitted that this use of the word "hyperplane" is not universal.]

A hyperplane need not be closed in E.

Proposition 11.3. *Let E be a real linear space, let p be a sublinear functional on E, and let F be a hyperplane in E. Suppose that $\phi : F \longrightarrow \mathbb{R}$ is a linear functional on F such that*

$$(\forall x \in F) \quad \phi(x) \le p(x)\,. \tag{42}$$

Then there is a linear functional $\psi : E \longrightarrow \mathbb{R}$ such that $\psi|F = \phi$ and

$$(\forall x \in E) \quad \psi(y) \le p(y)\,.$$

"If a linear functional on F is dominated by the sublinear functional p on E, then it may be extended to a linear functional on E that is also dominated by p."

Proof. Choose any $y_0 \in E \setminus F$; then, as F is a hyperplane in E, any $y \in E$ may be expressed in the form $y = x + \lambda y_0$ for some uniquely determined $\lambda \in \mathbb{R}$ and $x \in F$.

Let $x_1, x_2 \in F$. Then, by hypothesis (42),

$$\begin{aligned}\phi(x_1) - \phi(x_2) = \phi(x_1 - x_2) &\le p(x_1 - x_2) = p((x_1 + y_0) - (x_2 + y_0)) \\ &\le p(x_1 + y_0) + p(-x_2 - y_0)\,,\end{aligned}$$

and consequently

$$-\phi(x_2) - p(-x_2 - y_0) \le -\phi(x_1) + p(x_1 + y_0)\,.$$

This is so for any $x_1, x_2 \in F$. Now hold x_1 fixed, and then

$$\sup\{-\phi(x_2) - p(-x_2 - y_0) : x_2 \in F\} \le -\phi(x_1) + p(x_1 + y_0)\,,$$

and in turn, since this holds for any $x_1 \in F$,

$$\sup\{-\phi(x_2) - p(-x_2 - y_0) : x_2 \in F\} \leq \inf\{-\phi(x_1) + p(x_1 + y_0) : x_1 \in F\}. \quad (43)$$

Choose any $\tau \in \mathbb{R}$ between the left-hand and right-hand sides of (43), and then define, for any $\lambda \in \mathbb{R}$ and $x \in F$,

$$\psi(x + \lambda y_0) := \lambda\tau + \phi(x).$$

ψ is a well-defined linear functional $E \longrightarrow \mathbb{R}$, and $\psi|F = \phi$. It remains to show that it is dominated by p. If $x \in F$, then $\psi(x) = \phi(x) \leq p(x)$ by (42). Suppose that $y = x + \lambda y_0$, where $x \in F$ and $\lambda \neq 0$. Then, by the choice of τ, and referring to (43),

$$-\phi(\lambda^{-1}x) - p(-\lambda^{-1}x - y_0) \leq \tau \leq -\phi(\lambda^{-1}x) + p(\lambda^{-1}x + y_0).$$

If $\lambda > 0$, the right-hand inequality gives on multiplication by λ

$$\begin{aligned} \lambda\tau \leq -\lambda\phi(\lambda^{-1}x) + \lambda p(\lambda^{-1}x + y_0) = -\phi(x) + p(x + \lambda y_0) \quad &\text{and} \\ \psi(y) = \psi(x + \lambda y_0) = \lambda\tau + \phi(x) \leq p(x + \lambda y_0). \end{aligned}$$

If $\lambda < 0$, multiply the left-hand inequality by $-\lambda$:

$$\begin{aligned} \lambda\phi(\lambda^{-1}x) + \lambda p(-\lambda^{-1}x - y_0) \leq -\lambda\tau \quad &\text{and} \\ \psi(y) = \lambda\tau + \phi(x) \leq (-\lambda)p(-\lambda^{-1}x - y_0) = p(x + \lambda y_0). \end{aligned}$$

Therefore $\psi(y) \leq p(y)$ in all cases. □

An easy induction allows a similar extension of ϕ when F is of finite codimension in E.

Theorem 11.4. (The Hahn-Banach theorem.) *Let p be a sublinear functional on the real vector space E, in which F is a vector subspace. If $\phi : F \longrightarrow \mathbb{R}$ is a linear functional on F such that $\phi(x) \leq p(x)$ for all $x \in F$, then there exists a linear functional $\psi : E \longrightarrow \mathbb{R}$ such that $\psi|F = \phi$ and $\psi(y) \leq p(y)$ for all $y \in E$.*

Proof. Let $\mathcal{F}$ denote the set of pairs (G, χ) in which G is a vector subspace of E including F, and $\chi : G \longrightarrow \mathbb{R}$ is a linear functional dominated by p: $(\forall x \in G)\ \chi(x) \leq p(x)$. Introduce a partial order in $\mathcal{F}$ by

$$\text{“}(G_1, \chi_1) \leq (G_2, \chi_2)\text{” means “}G_1 \subseteq G_2 \text{ and } \chi_2|G_1 = \chi_1\text{”}. \quad (44)$$

That this is a partial order on $\mathcal{F}$ is trivial.

It is also inductive (recall §4); for suppose that $\{(G_\alpha, \chi_\alpha) : \alpha \in A\}$ is a chain in $\mathcal{F}$. Take $G_A := \bigcup_{\alpha\in A} G_\alpha$. If $x, y \in G$ and $\lambda, \mu \in \mathbb{R}$, there are indices $\beta, \gamma \in A$ such that $x \in G_\beta$ and $y \in G_\gamma$, but, as $\{(G_\alpha, \chi_\alpha)\}$ is a $\leq$-chain, either $G_\beta \subseteq G_\gamma$ or $G_\gamma \subseteq G_\beta$. Suppose, for instance, that $G_\beta \subseteq G_\gamma$. Then $x, y \in G_\gamma$, and so $\lambda x + \mu y \in G_\gamma \subseteq G$. This proves that G_A is a vector subspace of E. Define $\chi_A : G_A \longrightarrow \mathbb{R}$ in the obvious way: if $x \in G_\alpha$, $\chi_A(x) := \chi_\alpha(x)$. This is a satisfactory definition of χ_A, since, if $x \in G_\beta$ as well, then, as before, one of G_α, G_β includes the other and (by (44)) $\chi_\alpha(x) = \chi_\beta(x)$. And it is clear that, for each $\alpha \in A$, $(G_\alpha, \chi_\alpha) \leq (G_A, \chi_A)$.

From Zorn's Lemma, $\mathcal{F}$ has a maximal element (G, χ). If $x_0 \in E \setminus G$, then apply 11.3, taking "E" to be $G + \mathbb{R}x_0$, "p" to be $p|G + \mathbb{R}x_0$, "F" to be G, and "ϕ" to be χ. This would yield an element of $\mathcal{F}$ larger than (G, χ), which is absurd. Hence $G = E$, and χ may to be taken to be ψ in the statement of the theorem. □

In this argument the essential inductive stage, raising the dimension of the domain by 1, is 11.3. The problem is to show that ψ may be defined step by step by such one-dimensional extensions, and Zorn's Lemma is a convenient method of doing so. (One could also use "transfinite induction", i.e. the well-ordering principle). But, as usual, the consequence is that ψ is not described explicitly. When the Axiom of Choice was still regarded with suspicion, a version of the Hahn-Banach theorem was current that only required countably many steps (it was developed by Bohnenblust and Sobczyk), but the normed spaces had to be restricted to have "countable character" (formally, to be *separable*); and it is easy to invent normed or Banach spaces that are too large for that. The general Hahn-Banach theorem, for "arbitrarily large" normed spaces, might naturally be expected to appeal to the Axiom of Choice. Oddly enough, however, the existence of an *unbounded* linear functional also has to established via the Axiom, for instance by using an algebraic basis. It should be added that, on the 'classical' Banach spaces, all the bounded linear functionals are known — the force of the theorem is in its generality.

The Hahn-Banach theorem was originally proved in essentially the form I have given. In some sources, such as Bourbaki and the little book by the Robertsons, there is a more abstract and geometrical formulation (due, I think, to Bourbaki, who was possibly Schwartz). It removes explicit mention of p, whose significance is obscure, but it adds no extra generality and indeed is just a translation of the above reasoning into purely geometrical argument.

The theorem as given so far concerns the real case. For the complex case, a seminorm has to be used instead of a sublinear functional. The crucial point of the real proof is (43), which enables us to choose a suitable τ for the value $\psi(x_0)$; an analogous argument in the complex case would require τ to be a common point of all the closed balls $C(-\phi(y); p(y+x_0))$ as y varies over F, and (generally speaking) there is none, although any two of the balls do have a common point. In fact the complex case was only proved after some years' delay.

Lemma 11.5. *Let E be a complex vector space, and let $\phi : E \longrightarrow \mathbb{C}$. Define, for each $x \in E$, $\psi(x) := \Re(\phi(x))$ and $\chi(x) := \Im(\phi(x))$; thus $\psi, \chi : E \longrightarrow \mathbb{R}$. Then ϕ is complex-linear if and only if both ψ and χ are real-linear (when E is regarded as a real vector space) and $\chi(x) = -\psi(ix)$ for every $x \in E$.*

Proof. If ϕ is complex-linear, ψ and χ are trivially real-linear, and, for any $x \in E$,

$$\psi(ix) + i\chi(ix) = \phi(ix) = i\phi(x) = i(\psi(x) + i\chi(x)) = -\chi(x) + i\psi(x),$$

from which, taking real parts, $\chi(x) = -\psi(ix)$ as stated. Conversely, if ψ, χ are real-linear and $\chi(x) = -\psi(ix)$ for every x, and if $\alpha, \beta \in \mathbb{R}$,

$$\begin{aligned}
\phi((\alpha + i\beta)x) &= \psi((\alpha + i\beta)x) + i\chi((\alpha + i\beta)x) \\
&= \alpha\psi(x) + \beta\psi(ix) + i(\alpha\chi(x) + \beta\chi(ix)) \\
&= \alpha\psi(x) - \beta\chi(x) + \alpha i\chi(x) - \beta i\psi(i^2x) \\
&= \alpha\psi(x) - \beta\chi(x) + \alpha i\chi(x) + \beta i\psi(x) \\
&= (\alpha + i\beta)(\psi(x) + i\chi(x)) = (\alpha + i\beta)\phi(x).
\end{aligned}$$

□

Theorem 11.6. (The Hahn-Banach theorem, complex case.) *Let p be a seminorm on the complex vector space E; let F be a complex vector subspace of E, and let $\phi : F \longrightarrow \mathbb{C}$ be a linear functional such that $|\phi(x)| \leq p(x)$ for all $x \in F$. Then there is a linear functional $\Phi : E \longrightarrow \mathbb{C}$ such that $\Phi|F = \phi$ and $|\Phi(y)| \leq p(y)$ for all $x \in E$.*

Proof. Let $\psi(x) := \Re(\phi(x))$ and $\chi(x) := \Im(\phi(x))$ for all $x \in F$. Then ψ is a real-linear real-valued functional on F that is dominated by p:

$$(\forall x \in F) \quad \psi(x) \leq |\phi(x)| \leq p(x)\,.$$

By the real Hahn-Banach theorem 11.4, there is a real-linear functional $\Psi : E \longrightarrow \mathbb{R}$ such that $\Psi|F = \psi$ and $\Psi(y) \leq p(y)$ for all $y \in E$. Define

$$(\forall y \in E) \quad \Phi(y) := \Psi(y) - i\Psi(iy)\,.$$

By 11.5, Φ is a complex-linear functional $E \longrightarrow \mathbb{C}$; furthermore, if $x \in F$,

$$\Phi(x) = \psi(x) - i\psi(ix) = \psi(x) + i\chi(x) = \phi(x)\,.$$

It remains to show that Φ is still dominated by p. Given $y \in E$, choose $\theta \in \mathbb{R}$ such that $e^{i\theta}\Phi(y)$ is real and positive. Then

$$|\Phi(y)| = e^{i\theta}\Phi(y) = \Phi(e^{i\theta}y) = \Psi(e^{i\theta}y) \leq p(e^{i\theta}y) = |e^{i\theta}|\, p(y) = p(y)\,. \qquad \square$$

§12. Some consequences of the Hahn-Banach theorem

Recall that the dual of a normed space E has a norm induced from the norm on E; it is the operator-norm in $L(E;\mathbb{K})$, but in this context may be called the *dual norm*.

Theorem 12.1. *Let F be a subspace of a normed space E, and let ϕ be a bounded linear functional on F. Then there is a bounded linear functional Φ on E such that $\Phi|F = \phi$ and $\|\Phi\|_{E'} = \|\phi\|_{F'}$.*

Proof. By 10.6, $|\phi(x)| \leq \|\phi\|_{F'}\|x\|_F = \|\phi\|_{F'}\|x\|_E$ for all $x \in F$. Define the seminorm p in 11.6 by $p(y) := \|\phi\|_{F'}\|y\|_E$ for all $y \in E$. Thus there is a linear functional $\Phi : E \longrightarrow \mathbb{K}$ such that $\Phi|F = \phi$ and $|\Phi(y)| \leq \|\phi\|_{F'}\|y\|_E$, so that $\|\Phi\|_{E'} \leq \|\phi\|_{F'}$. However, it is trivially obvious that $\|\Phi\|_{E'} \geq \|\phi\|_{F'}$. $\square$

Corollary 12.2. *Let E be a normed space, and suppose $0 \neq x_0 \in E$. Then there exists $\Phi \in E'$ such that $\|\Phi\|_{E'} = 1$ and $\Phi(x_0) = \|x_0\|_E$.*

Proof. Take $F := \{\alpha x_0 : \alpha \in \mathbb{K}\}$ and $\phi(\alpha x_0) := \alpha\|x_0\|_E$ for all $\alpha \in \mathbb{K}$. Then $\phi \in F'$ and $\|\phi\|_{F'} = 1$. Apply the Theorem. $\square$

One might ask whether there is a "reverse" to this result, saying that, for any $\phi \in E'$, there is some $x \in E$ such that $\|x\|_E = 1$ and $\phi(x) = \|\phi\|_{E'}$. This need not be true. (See exercise set 4, no. 3).

Remark 12.3. To the more general question whether, given a vector subspace F of E and another normed space G (of dimension higher than 1), any bounded linear map $S : F \longrightarrow G$ has a bounded linear extension $T : E \longrightarrow G$, there is no satisfactory general answer. If F is finite-dimensional, with a basis $g_1, g_2, \ldots, g_n$, then there are $\phi_1, \phi_2, \ldots, \phi_n \in F'$ such that

$$(\forall x \in F) \quad Sx = \phi_1(x)g_1 + \phi_2(x)g_2 + \cdots + \phi_n(x)g_n\,.$$

(That the ϕ_k are all continuous is not obvious; just assume it for the moment.) Then, each ϕ_k has a bounded extension $\Phi_k \in E'$, and S may be extended to $T : E \longrightarrow G$, where

$$(\forall y \in E) \quad Ty := \Phi_1(x)g_1 + \Phi_2(x)g_2 + \cdots + \Phi_n(x)g_n .$$

However, $\|T\|_{L(E,G)}$ may well be larger than $\|S\|_{L(F,G)}$.

Of course, whether G is finite-dimensional or not, there is a *linear* extension of S to a linear mapping $E \longrightarrow G$; one needs only to extend an algebraic basis of F to a basis of E (see 7.1). But this extension may well be unbounded.

Corollary 12.4. *Let E, G be normed spaces and $0 \neq x_0 \in E$, $y_0 \in G$. Then there exists $T \in L(E,G)$ such that $\|T\|_{L(E,G)} = \|y_0\|/\|x_0\|$ and $Tx_0 = y_0$.*

Proof. Let $Tx := \Phi(x)y_0$, where Φ is as in 12.2. □

This resolves the question whether $L(E,G)$ has non-trivial members.

Definition 12.5. Let E be a normed space. The *bidual* of E is the (continuous) dual E'' of the (continuous) dual E' of E, with the operator-norm. The *bidual map* $J : E \longrightarrow E''$ is defined by setting, for any $x \in E$, $(Jx)\phi := \phi(x)$ for each $\phi \in E'$. Compare 8.3; as there, we may write $\langle \phi, Jx \rangle = \langle x, \phi \rangle$.

J is linear, and, for any $\phi \in E'$, $|(Jx)\phi| = |\phi(x)| \leq \|\phi\|\|x\|$, so $\|Jx\| \leq \|x\|$.

Theorem 12.6. *If E is a normed space, the bidual map $J : E \longrightarrow E''$ is an isometry (and, consequently, injective).*

Proof. Given $0 \neq x \in E$, there is $\phi \in E'$ such that $\|\phi\| = 1$ and $\phi(x) = \|x\|$, by 12.2. But then $\|Jx\|_{E''} \geq |\langle \phi, Jx \rangle| = |\langle x, \phi \rangle| = \|x\|_E$, and we have already seen $\|Jx\| \leq \|x\|$.□

Recall from 8.4 that the “algebraic bidual map” $E \longrightarrow E^{**}$ is injective too. But it may be surjective when E is infinite-dimensional, unlike its algebraic analogue.

Definition 12.7. The normed space E is *reflexive* if the bidual map J is surjective; that is, if every continuous linear functional on E' is of the form $\phi \mapsto \phi(x)$ for some $x \in E$.

In this case, 12.6 ensures that J is an isometric isomorphism between E and E''. This can only occur if E is complete, for E'' is automatically complete by 10.15. But amongst Banach spaces, many are non-reflexive and a significant number are reflexive. For instance, any Hilbert space is reflexive (as we shall soon see), as are all the spaces L^p for $1 < p < \infty$; but the space c_0 of $\mathbb{K}$-valued sequences convergent to 0, with supremum norm, is not reflexive. Its dual may be identified with l^1, whose dual in turn may be identified with l^∞.

Definition 12.8. Let E be a normed space and $A \subseteq E$. A is *norm-bounded* (or simply *bounded*) if there is a number $K \geq 0$ such that $\|x\| \leq K$ for all $x \in A$. A is *weakly bounded* (or *$\sigma(E,E')$-bounded*) if, for any $\phi \in E'$, there is $K_\phi \geq 0$ such that $|\phi(x)| \leq K_\phi$ for all $x \in A$.

The *polar set* of A is the set $A^0 \subseteq E'$ defined by

$$A^0 := \{\phi \in E' : (\forall x \in A) \ \ |\phi(x)| \leq 1\}.$$

It is easy to see that A^0 is convex, balanced, and closed (in the norm-topology) in E'.

Theorem 12.9. *Let A be a weakly bounded subset of the normed space E. Then A is also norm-bounded.*

Proof. For each $\phi \in E'$, there is an $n \in \mathbb{N}$ such that $|\phi(x)| \leq n$ for all $x \in A$; thus $\phi \in nA^0$. Hence, $E' = \bigcup_{n=1}^{\infty} nA^0$. But E' is a Banach space by 10.15, and so of the second Baire category; for some n, the closed set nA^0 must have non-empty interior. Thus A^0 also has non-empty interior, and includes $B(\phi_0; 2\epsilon)$ for some $\phi_0 \in E'$ and $\epsilon > 0$. As it is balanced, $A^0 \supseteq B(-\phi_0; 2\epsilon)$ too. As it is convex, $B(0; 2\epsilon) \subseteq A^0$. (Indeed, if $\phi \in B(0; 2\epsilon)$, $\phi = \frac{1}{2}(\phi + \phi_0) + \frac{1}{2}(\phi - \phi_0)$, where $\phi + \phi_0 \in B(\phi_0; 2\epsilon)$ and $\phi - \phi_0 \in B(-\phi_0; 2\epsilon)$.)

Thus, whenever $\phi \in E'$ and $\|\phi\|_{E'} \leq \epsilon$, $|\phi(x)| \leq 1$ for all $x \in A$. (*)

Suppose, if possible, that $x \in A$ and $\|x\|_E > \epsilon^{-1}$. By 12.2, there is $\Phi \in E'$ such that $\|\Phi\|_{E'} = 1$ and $\Phi(x) = \|x\|_E > \epsilon^{-1}$. Take $\phi := \epsilon\Phi$, and then $\|\phi\| = \epsilon$ and $|\phi(x)| > 1$. This contradicts (*), and we must conclude that $A \subseteq C(0; \epsilon^{-1})$. □

This result is sometimes called the *Principle of Uniform Boundedness*. Its proof may be compared with the rather similar 9.13 (which, in Banach spaces, is sometimes called by the same name). There is an alternative argument, which I shall now sketch.

Consider the space $\mathcal{F}$ of *all* bounded functions $A \longrightarrow \mathbb{K}$, with supremum norm. $\mathcal{F}$ is a Banach space, and the weak boundedness of A implies that there is a mapping $R : E' \longrightarrow \mathcal{F}$ defined by restricting $\phi \in E'$ to A. The graph $G(R)$ of R is closed (as is easily checked). By the closed graph theorem 9.11, R is continuous, $R \in L(E', \mathcal{F})$; and $|\phi(x)| \leq \|R\|\|\phi\|$ for all $x \in A$ and all $\phi \in E'$. From 12.2, this suffices to show that $\|x\| \leq \|R\|$ for all $x \in A$.

Remark 12.10. If E, F are normed spaces and $T \in L(E, F)$, there is an induced mapping $T' \in L(F', E')$ (the *dual mapping* or *conjugate mapping*) defined by

$$(\forall \phi \in F')(\forall x \in E) \quad \langle x, T'\phi \rangle := \langle Tx, \phi \rangle$$

(the first dual pairing is between E and E', the second between F and F'). Equivalently, $T'\phi : E \longrightarrow \mathbb{K}$ is $\phi \circ T$, which shows at once that $T'\phi \in E'$. In fact, $\|T'\| = \|T\|$.

The correspondence from E to E' is a *contravariant functor* from the category of normed spaces and bounded linear maps to itself. In particular, $(ST)' = T'S'$ whenever possible. And $T''J = JT$ (which, in effect, means that T'' extends T from E to the larger space E''). J is a *natural transformation* between the identity functor and the "bidual" functor.

Perhaps the most surprising application of the Hahn-Banach theorem is the construction of "generalized limits". Consider the space l^∞ of bounded $\mathbb{K}$-valued sequences, with norm

$$\|(\xi_n)\| := \sup\{|\xi_n| : n \in \mathbb{N}\}.$$

There is a (closed) vector subspace (usually called c) consisting of all the *convergent* sequences, and the limit is a linear functional on c dominated by the norm:

$$(\forall (\xi_n) \in c) \quad |\lim_{n\to\infty} \xi_n| \leq \|(\xi_n)\|.$$

Hence, by the Hahn-Banach theorem, there is a linear functional $\phi : l^\infty \longrightarrow \mathbb{K}$ such that

$$(\forall(\xi_n) \in c) \;\; \phi((\xi_n)) = \lim_{n\to\infty} \xi_n\,, \quad \text{and} \quad (\forall(\xi_n) \in l^\infty) \;\; |\phi((\xi_n))| \leq \|(\xi_n)\|\,.$$

In this sense, one may assign to each bounded sequence a "limit". But, apart from agreeing with the ordinary limit when that exists, ϕ need not have any limit-like properties. Banach and Mazur pointed out that, by an ingenious choice of sublinear functional "p", one can force any linear functional "ϕ" dominated by p (not necessarily constructed by extending the ordinary limit on c) to behave in some respects quite like a limit. Of course, since the construction involves Zorn's lemma, ϕ cannot be described explicitly (and is not, in fact, unique). The trick has some other striking applications, which, for brevity, I omit; you will find them in Banach's book *Théorie des opérations linéaires*, ch. II, §3.

§13 Hilbert spaces: the dual space.

In this section, H denotes a Hilbert space over $\mathbb{K}$ with inner product $\langle,\rangle$ and induced norm $\|\|$. If $a, b \in H$, it is often convenient to write $a \perp b$ to mean $\langle a, b\rangle = 0$; the relation $\perp$ is symmetric. If $B \subseteq H$, $a \perp B$ (or $B \perp a$) means that $a \perp b$ for all $b \in B$. Similarly, if $A, B \subseteq H$, $A \perp B$ (or $B \perp A$) means $a \perp b$ for all $a \in A$ and $b \in B$.

Lemma 13.1. *If* $a \perp b$, *then* $\|a+b\|^2 = \|a\|^2 + \|b\|^2$.

Proof. $\langle a+b, a+b\rangle = \langle a, a\rangle + \langle a, b\rangle + \langle b, a\rangle + \langle b, b\rangle = \langle a, a\rangle + \langle b, b\rangle$. □

This is of course the "Theorem of Pythagoras".

Lemma 13.2. *Let* $x, y \in H$. *Then* $\|x+y\|^2 + \|x-y\|^2 = 2\|x\|^2 + 2\|y\|^2$.

This is of course the classical "Theorem of Apollonius".

Proof.
$$\begin{aligned}\|x+y\|^2 + \|x-y\|^2 &= \langle x+y, x+y\rangle + \langle x-y, x-y\rangle \\ &= \langle x, x\rangle + \langle x, y\rangle + \langle y, x\rangle + \langle y, y\rangle + \langle x, x\rangle - \langle x, y\rangle - \langle y, x\rangle + \langle y, y\rangle \\ &= 2\langle x, x\rangle + 2\langle y, y\rangle = 2\|x\|^2 + 2\|y\|^2\,.\end{aligned}$$ □

Lemma 13.3. *Let K be a non-null closed convex set in H. For any* $a \in H$, *there is a unique* $z \in K$ *such that* $\|a - z\| = d(a, K) := \inf\{\|a-y\| : y \in K\}$.

Proof. For each $m \in \mathbb{N}$, there is $z_m \in K$ such that $\|a - z_m\| < \sqrt{d(a,K)^2 + m^{-1}}$. Apply 13.2 with $y := \frac{1}{2}(z_m - z_n)$ and $x := a - \frac{1}{2}(z_m + z_n)$; then $\frac{1}{2}(z_m + z_n) \in K$ too, as K is convex, and so $\left\|a - \frac{1}{2}(z_m + z_n\right\|^2 \geq d(a,K)^2$:

$$\begin{aligned}2\|z_m - z_n\|^2 &= \|a - z_n\|^2 + \|x - z_m\|^2 - 2\left\|a - \tfrac{1}{2}(z_m + z_n)\right\|^2 \\ &< d(a,K)^2 + n^{-1} + d(a,K)^2 + m^{-1} - 2d(a,K)^2 \\ &= m^{-1} + n^{-1}\,.\end{aligned}$$

This shows that (z_m) is Cauchy, and so has a limit z in H, which must belong to K. And

$$d(a,K) \le \|a-z\| \le \|a-z_n\| + \|z_n - z\| < \sqrt{d(a,K)^2 + n^{-1}} + \|z_n - z\|$$

for all n, which, since $\|z_n - z\| \to 0$ as $n \to \infty$, implies that $\|a-z\| = d(a,K)$.

If $z' \in K$ also satisfies this equality, then

$$\begin{aligned} 0 \le 2\|z-z'\|^2 &= \|a-z\|^2 + \|a-z'\|^2 - 2\big\|a - \tfrac{1}{2}(z+z')\big\|^2 \\ &\le d(a,K)^2 + d(a,K)^2 - 2d(a,K)^2 = 0\,, \end{aligned}$$

and it follows that $z = z'$. □

The proof uses only the completeness of K, not of H. From an applied mathematician's viewpoint, the result says that K must contain a "closest approximation" to a.

Lemma 13.4. *Let K be a proper closed vector subspace of H, and $a \notin K$. If $z \in K$, then $d(a,K) = \|a-z\|$ if and only if $a - z \perp K$.*

Proof. Suppose $d(a,K) = \|a-z\|$, and let $y \in K$. Then, for any $\lambda \in \mathbb{K}$,

$$\begin{aligned} \|a-z\|^2 \le \|a-z+\lambda y\|^2 &= \langle a-z+\lambda y, a-z+\lambda y\rangle \\ &= \langle a-z, a-z\rangle + \lambda\langle y, a-z\rangle + \overline{\lambda}\langle a-z, y\rangle + \lambda\overline{\lambda}\langle y, y\rangle\ . \end{aligned}$$

Take $\lambda := -\rho\langle a-z, y\rangle$, where ρ is real and positive; then, cancelling $\|a-z\|^2$,

$$0 \le -2\rho|\langle a-z, y\rangle|^2 + \rho^2|\langle a-z, y\rangle|^2\|y\|^2\,.$$

If $\langle a-z, y\rangle \ne 0$, this is false when $\rho\|y\|^2 < 2$. Hence $a - z \perp y$.

Conversely, let $a - z \perp K$. Given any $y \in K$, $z - y \in K$ too, $a - z \perp z - y$, and (by 13.1) $\|a-y\|^2 = \|y-z\|^2 \ge \|a-z\|^2$; hence, indeed, $\|a-z\| = d(a,K)$. □

Corollary 13.5. *If K is a proper closed subspace of H, there exists some $b_2 \in H$ such that $0 \ne b_2 \perp K$.* (Notice that this implies $b_2 \notin K$.)

Proof. As K is proper, one may take $a \notin K$, and then construct $b_2 := a - z$ as above. □

Lemma 13.6. *If K is a closed non-trivial subspace of H, define $P_K : H \longrightarrow K$ by*

$$(\forall a \in H) \quad P_K a := z\,,$$

where z is the unique element of K such that $a - z \perp K$. (See 13.3 *and* 13.4.*) Then P_K is a bounded linear mapping, and $\|P_K\| = 1$.*

Proof. If $\lambda, \mu \in \mathbb{K}$ and $a, b \in H$, then $a - P_K a \in K^\perp$ and $b - P_K b \in K^\perp$, and so

$$(\lambda a + \mu b) - (\lambda P_K a + \mu P_K b) = \lambda(a - P_K a) + \mu(b - P_K b) \in K^\perp\,,$$

which (by uniqueness in 13.3) implies that $P_K(\lambda a + \mu b) = \lambda P_K a + \mu P_K b$. Furthermore, as

$$a = P_K a + (a - P_K a) \quad \text{and} \quad P_K a \perp a - P_K a\,,$$

13.1 shows that $\|a\|^2 = \|P_K a\|^2 + \|a - P_K a\|^2 \geq \|P_K a\|^2$. Hence, $\|P_K\| \leq 1$. However, K is non-trivial, so there is some non-zero $k \in K$, and $P_K k = k$. Thus $\|P_K\| = 1$. ☐

P_K is usually called "orthogonal projection on K". If K is the zero subspace, then, of course, P_K is the zero mapping. At 17.9 we shall have an "explicit" description of P_K.

Theorem 13.7. *There is a mapping* $Q : H \longrightarrow H'$ *defined by*

$$(\forall x, y \in H) \quad Q(y)(x) = \langle x, y \rangle .$$

It is a bijective correspondence, isometric, and conjugate-linear.

Proof. $Q(y)$ is a linear functional because $\langle , \rangle$ is linear in the first variable. It is a *bounded* linear functional by the Cauchy-Schwarz inequality: $|Q(y)(x)| = |\langle x, y \rangle| \leq \|x\|\|y\|$. Moreover, this shows that $\|Q(y)\|_{H'} \leq \|y\|$. On the other hand,

$$Q(y)y = \langle y, y \rangle = \|y\|^2 ,$$

which implies, if $y \neq 0$, that $\|Q(y)\|_{H'} \geq \|y\|$. Hence, $\|Q(y)\| = \|y\|$ when $y \neq 0$; and this is also true when $y = 0$. Q is, therefore, isometric (and so injective). It is conjugate-linear because $\langle , \rangle$ is conjugate-linear in the second variable.

It remains to prove Q *surjective*. Let $0 \neq \phi \in H'$, and $K := \phi^{-1}(\{0\})$. K is a proper closed vector subspace of H. By 13.5, there is some $b_2 \notin K$ such that $b_2 \perp K$; thus, $\phi(b_2) \neq 0$, and I take $b_1 := \phi(b_2)^{-1} b_2$.

Given $x \in H$, $\phi(x) = \phi(\phi(x) b_1)$, and so $x - \phi(x) b_1 \in K$. But, as $b_1 \perp K$,

$$0 = \langle x - \phi(x) b_1, b_1 \rangle = \langle x, b_1 \rangle - \phi(x) \langle b_1, b_1 \rangle , \quad \phi(x) = \langle x, \langle b_1, b_1 \rangle^{-1} b_1 \rangle .$$

Consequently, if $b := \langle b_1, b_1 \rangle^{-1} b_1$, then $\phi(x) = \langle x, b \rangle$ for all $x \in H$, and $\phi = Q(b)$.

If $\phi = 0 \in H'$, it is clear that $\phi = Q(0_H)$. ☐

Thus we may identify a Hilbert space with its own dual, except for the inconvenient fact that Q is a *conjugate*-isometry. In the real case, it is literally true that Q is an isomorphism of H with H'. In general, we turn H' into a Hilbert space by defining, in an obvious notation,

$$\langle Q_H(x), Q_H(y) \rangle_{H'} := \langle y, x \rangle_H ,$$

and then $(Q_{H'} Q_H(y))(Q_H(x)) = \langle y, x \rangle_H = (Q_H(x))y$, which means $J(y) = Q_{H'} Q_H(y)$.

Corollary 13.8. *Any Hilbert space is reflexive.* ☐

Indeed, the bidual map J is the composite of the bijections $Q_{H'}$ and Q_H.

It should be noted here that I have used the notation $\langle , \rangle$ both for the inner product in H and for the dual pairing between E and E^* (see 8.3). There is a difference, since the inner product is not linear but conjugate-linear in the second argument. In the real case, we can regard H as its own dual, paired with itself by the inner product; but in the complex case, H', given a Hilbert space structure as above, is in a sense the "complex conjugate" of H.

Theorem 13.7 is known in some quarters as the *Riesz representation theorem*, since it gives an exact description of the continuous linear functionals on H. I am not sure of the history, but the phrase "Riesz representation theorem" is associated in most pure mathematicians' minds with a far more substantial theorem about bounded linear functionals on the Banach spaces $C(\Omega;\mathbb{R})$ of continuous functions.

§14. The Radon-Nikodym theorem. (2008: omit.)

§15. Duality in L^p spaces. (2008: omit.)

§16. Conditional expectation. (2008: omit.)

§17. Hilbert spaces: classification.

Throughout this section H is a Hilbert space. (Some of the definitions and results require only that it should be an inner product space.)

Definition 17.1. If $A \subseteq H$, let $A^{\perp} := \{y \in H : (\forall x \in A)\ x \perp y\}$.

Lemma 17.2. *For any $A \subseteq H$, $A^{\perp}$ is a closed vector subspace of H.* ☐

The proof that $A^{\perp}$ is closed appeals to the Cauchy-Schwarz inequality. It will be convenient to use the notation $\Lambda(A)$ to denote cl(Span(S)).

Definition 17.3. A subset S of H is *orthonormal* if, whenever $a, b \in S$,

$$\langle a, b\rangle = \begin{cases} 1 & \text{if } a = b, \\ 0 & \text{if } a \neq b. \end{cases}$$

An orthonormal set S in H is described as *complete* or *fundamental* if $\Lambda(S) = H$.

In the next few results S denotes an orthonormal set in H.

Lemma 17.4. (Bessel's inequality.) *Let S be finite. For any $x \in H$,*

$$\|x\|^2 \geq \sum_{a\in S} |\langle x, a\rangle|^2, \tag{45}$$

with equality if and only if $x \in$ Span(S).

Proof.
$$\begin{aligned}0 &\le \Big\|x - \sum_{a\in S} \langle x,a\rangle a\Big\|^2 = \Big\langle x - \sum_{a\in S}\langle x,a\rangle a,\, x - \sum_{b\in S}\langle x,b\rangle b\Big\rangle \\
&= \langle x,x\rangle - \Big\langle x, \sum_{b\in S}\langle x,b\rangle b\Big\rangle - \Big\langle \sum_{a\in S}\langle x,a\rangle a,\, x\Big\rangle \\
&\qquad + \Big\langle \sum_{a\in S}\langle x,a\rangle a,\, \sum_{b\in S}\langle x,b\rangle b\Big\rangle \\
&= \langle x,x\rangle - \sum_{b\in S}\overline{\langle x,b\rangle}\langle x,b\rangle - \sum_{a\in S}\langle x,a\rangle\langle a,x\rangle \\
&\qquad + \sum_{a,b\in S}\langle x,a\rangle\overline{\langle x,b\rangle}\langle a,b\rangle \qquad (46)\\
&= \langle x,x\rangle - \sum_{b\in S}|\langle x,b\rangle|^2 - \sum_{a\in S}|\langle x,a\rangle|^2 + \sum_{a\in S}|\langle x,a\rangle|^2 \qquad (47)\\
&= \langle x,x\rangle - \sum_{a\in S}|\langle x,a\rangle|^2 .\end{aligned}$$

Notice that the passage from (46) to (47) uses the fact that $\langle a,b\rangle = 0$ unless $a = b$. This proves (45). If it is an equality, then $\|x - \sum_{a\in S}\langle x,a\rangle a\| = 0$, and so

$$x = \sum_{a\in S}\langle x,a\rangle a \in \text{Span}(S)\,.$$

Finally, suppose that $x = \sum_{a\in S}\lambda_a a \in \text{Span}(S)$. Then, for each $b \in S$,

$$\langle x,b\rangle = \sum_{a\in S}\lambda_a\langle a,b\rangle = \lambda_b\,,$$

so that, in fact, $x = \sum_{a\in S}\langle x,a\rangle a$, and the calculation above shows (45) is an equality. □

Remark 17.5. I recall from 441 (where, indeed, I did not really expound the ideas in any detail) that, in any topological vector space E, an indexed subset $\{v_\alpha : \alpha\in A\}$ is said to have an *unordered sum* $\sigma \in E$ if, for any open set V containing σ, there is a finite subset F of A with the property that, for any finite subset G of A such that $F \subseteq G$, $\sum_{\alpha\in G} v_\alpha \in V$.

If E is a Banach space, $\{v_\alpha : \alpha\in A\}$ has an unordered sum if and only if it satisfies the following Cauchy condition: for any $\epsilon > 0$, there exists a finite subset F of A such that, for any finite subset G of A disjoint from F, $\|\sum_{\alpha\in G} v_\alpha\| < \epsilon$.

Corollary 17.6. *Let S be* any *orthonormal set. Then, for any $x\in H$, the S-indexed set of non-negative real numbers $\{|\langle x,a\rangle|^2 : a\in S\}$ has an unordered sum $\sum_{a\in S}|\langle x,a\rangle|^2$, and*

$$\|x\|^2 \ge \sum_{a\in S}|\langle x,a\rangle|^2 .$$

Proof. The sums over finite subsets of the index set S are all bounded above by $\|x\|^2$. The unordered sum is (all terms being non-negative) the supremum of the finite sums. □

Corollary 17.7. *For any orthonormal set S in H and any $x\in H$, the subset*

$$S[x] := \{a\in S : \langle x,a\rangle \ne 0\}$$

is countable.

Proof. For $n\in\mathbb{N}$, let $S[x,n] := \{a\in S : |\langle x,a\rangle| \ge 1/n\}$. Suppose F is a finite subset of $S[x,n]$. Then 17.4 shows $\|x\|^2 \ge \sum_{a\in F}|\langle x,a\rangle|^2$, and as each term of the sum is at least n^{-2}, it follows that $\#(F) \le n^2\|x\|^2$. Since this is true for any finite subset of $S[x,n]$, necessarily $S[x,n]$ itself is finite with at most $n^2\|x\|^2$ elements. Consequently,

$$S[x] = \bigcup_{n\in\mathbb{N}} S[x,n]$$

must be countable. $\square$

This fact enables textbooks to avoid the explicit notion of an unordered sum (and not only in this context).

Lemma 17.8. *Let S be finite and $x \in H$. Then $\|\sum_{a\in S} \langle x,a\rangle a\|^2 = \sum_{a\in S} |\langle x,a\rangle|^2$.* $\square$

Corollary 17.9. *For any orthonormal S and any $x \in H$, the unordered sum $\sum_{a\in S} \langle x,a\rangle a$ exists and is in the closure of* $\mathrm{Span}(S)$.

Proof. Given $\epsilon > 0$, there is a finite subset F of S such that

$$\sum_{a\in F} |\langle x,a\rangle|^2 > \sum_{a\in S} |\langle x,a\rangle|^2 - \epsilon^2,$$

and, for any finite subset G of S disjoint from F, it must, therefore, be the case that

$$\left\|\sum_{a\in G} \langle x,a\rangle a\right\|^2 = \sum_{a\in G} |\langle x,a\rangle|^2 < \epsilon^2. \tag{48}$$

This means that the Cauchy condition of 17.5 is satisfied, so the unordered sum $\sigma := \sum_{a\in S} \langle x,a\rangle a$ exists. But then, again given any $\epsilon > 0$, there is some finite subset F of S [indeed, for the same ϵ as before one may take the same F as before, but it is not necessary to prove this here] such that $\|\sigma - \sum_{a\in F} \langle x,a\rangle a\| < \epsilon$, and $\sum_{a\in F} \langle x,a\rangle a \in \mathrm{Span}(S)$; so, as asserted, $\sigma \in \mathrm{cl}(\mathrm{Span}(S))$.

Lemma 17.10. *Let $\{y_\alpha : \alpha \in A\}$ be any indexed subset of H such that the unordered sum $\sum_{\alpha\in A} y_\alpha$ exists in H and is equal to $y \in H$. Then, for any $z \in H$, the unordered sum $\sum_{\alpha\in A} \langle y_\alpha, z\rangle$ exists in $\mathbb{K}$ and is equal to $\langle y,z\rangle$.*

Proof. Given $\epsilon > 0$, there exists a finite subset F of A such that, for any finite subset G of A that includes F, $\|y - \sum_{\alpha\in G} y_\alpha\| < (1+\|z\|)^{-1}\epsilon$. By the Cauchy-Schwarz inequality,

$$\begin{aligned}\left|\langle y,z\rangle - \left\langle \sum_{\alpha\in G} y_\alpha, z\right\rangle\right| &= \left|\left\langle y - \sum_{\alpha\in G} y_\alpha, z\right\rangle\right| \\ &\le \left\|y - \sum_{\alpha\in G} y_\alpha\right\| \|z\| \le (1+\|z\|)^{-1}\epsilon\|z\| < \epsilon.\end{aligned}$$

This establishes the Lemma. $\square$

Lemma 17.11. *Given any orthonormal set S in H, let $\Lambda(S) := \mathrm{cl}(\mathrm{Span}(S))$, and let $P_{\Lambda(S)}$ be the orthogonal projection on $\Lambda(S)$ (see* 13.6*). Then, for any $x \in H$,*

$$\sum_{a\in S} \langle x,a\rangle a = P_{\Lambda(S)} x,$$

where the sum exists by 17.9.

Proof. Indeed, the sum is in $\Lambda(S)$ by 17.9, and, for any $b \in S$, by 17.10

$$\left\langle x - \sum_{a\in S} \langle x,a\rangle a, b \right\rangle = \langle x,b\rangle - \sum_{a\in S} \langle x,a\rangle\langle a,b\rangle = \langle x,b\rangle - \langle x,b\rangle = 0,$$

so that $S \subseteq \{x - \sum_{a\in S} \langle x,a\rangle a\}^{\perp}$, and $\Lambda(S) \subseteq \{x - \sum_{a\in S} \langle x,a\rangle a\}^{\perp}$ by 17.2. $\square$

Remark 17.12. The familiar "Gram-Schmidt process" is an application of this argument. If $0 \neq c \notin \Lambda(S)$, then $0 \neq d = c - \sum_{a\in S} \langle c,a\rangle a \perp \Lambda(S)$, and $S \cup \{\|d\|^{-1} d\}$ is an orthonormal set which includes S; indeed, it has the same closed span as $S \cup \{c\}$. It follows that

Lemma 17.13. *An orthonormal set S in H is complete if and only if it is maximal in the class of orthonormal sets ordered by inclusion.*

Proof. If S is not complete, $\Lambda(S) \neq H$, and the Remark provides an orthonormal set strictly including S. On the other hand, if S is not maximal, there is a vector $c \in S$ such that $S \cup \{c\}$ is orthonormal, and then $S \subseteq \{c\}^{\perp}$ and $\Lambda(S) \subseteq \{c\}^{\perp}$ by 17.2. Thus $c \notin \Lambda(S)$, and S is not complete. $\square$

Proposition 17.14. *Let S_0 be any orthonormal set in H. There exists a complete orthonormal set S that includes S_0.*

Proof. Consider the class $\mathcal{S}$ of all orthonormal sets that include S_0, and order it by inclusion. It is clear that the union of any chain in $\mathcal{S}$ is also an orthonormal set. By Zorn's Lemma there is a maximal element in $\mathcal{S}$ (and, therefore, in the class of *all* orthonormal sets ordered by inclusion). But 17.13 shows that this maximal orthonormal set is complete.

Remark 17.15. A complete orthonormal set S in H is often called an *orthonormal basis* (or, sometimes, a *Hilbert basis*) in H. The name "basis" is not entirely inappropriate, since it follows from the results above that any vector $x \in H$ may be expressed as an unordered sum $\sum_{a\in S} \langle x,a\rangle a$, and that this sum may in turn be expressed as the unconditionally convergent (i.e. unchanged under arbitrary rearrangement) sum of a series $\sum_{n=1}^{\infty} \langle x,a_n\rangle a_n$ for some sequence (a_n) of elements of S (the sequence depends on x). Furthermore, the coefficients have what is sometimes called "finality": if we take any finite subset F of S, then $\sum_{a\in F} \langle x,a\rangle a$ is the best approximation to x in $\mathrm{Span}(F)$; the coefficient of a in this best approximation is the same for any F that contains a.

Any vector of the space is a (finite) linear combination of elements of an *algebraic* basis. The existence of an algebraic basis was established at 7.1 by Zorn's Lemma. An "orthonormal basis" is not an algebraic basis (unless it is itself finite); but it too exists, in general, only because of the Axiom of Choice. There are, however, well-known explicit examples of orthonormal bases in interesting spaces.

Proposition 17.16. *In the space $L^2((0,2\pi);\lambda)$ of square-integrable complex-valued Lebesgue-measurable functions on the interval $(0,2\pi)$, the set $\left\{\dfrac{\exp(int)}{\sqrt{2\pi}} : n \in \mathbb{Z}\right\}$ is an orthonormal basis.* $\square$

"Proof". It is easy to check that the set is orthonormal, so the only problem is to show that it is complete. Since the product of two functions in the set is another, its linear span is a subalgebra A of the algebra C of 2π-periodic continuous functions on $\mathbb{R}$, which is identified with $C(\mathbb{T};\mathbb{C})$ (compare the argument after 6.13). In this sense A separates the points of $\mathbb{T}$, and contains the constant functions; so, by the Stone-Weierstrass theorem, it is dense in the space of all continuous functions $\mathbb{T} \longrightarrow \mathbb{C}$ in the supremum norm, which is in effect C.

Given $f \in L^2((0,2\pi);\lambda)$ and $\epsilon > 0$, there is some function $g \in C$ such that $\|f-g\|_{L^2} < \frac{1}{2}\epsilon$. (This is true; but I have not proved it, here or in 441.) By the preceding remarks, there exists $h \in A$ such that $\|h-g\|_\infty < \epsilon/\sqrt{8\pi}$, where $\|.\|_\infty$ denotes the supremum norm, and then $\|h-g\|_{L^2} < \frac{1}{2}\epsilon$ and $\|f-h\| < \epsilon$. $\square$

There are other proofs of this result, depending on integration by parts, which I have also not discussed. But many of the results I have expounded, such as Bessel's inequality, were originally proved for the basis above, or the Fourier basis which may be deduced from it, and for continuous functions. In fact, Lebesgue's theory of integration was perhaps motivated by the need to make more sense of results on Fourier series. It may be added that the coefficients $\langle x, a\rangle$ (for $a \in S$) are sometimes called the *Fourier coefficients* of x with respect to the orthonormal set S.

Theorem 17.17. *Any two orthonormal bases of H have the same cardinal number.*

Proof. If H is finite-dimensional (in the algebraic sense), an orthonormal basis is an algebraic basis (why?) and the result is standard. If H is infinite-dimensional, any orthonormal basis must be infinite. Let S_1 and S_2 be two orthonormal bases, with infinite cardinalities $|S_1|$, $|S_2|$. Let X_1 be the set of rational-complex (or rational-real in the real case) finite linear combinations of elements of S_1. Then X_1 is dense in H (see 17.15; the partial sums of the series approximate x and may themselves be approximated by rational linear combinations). However, $|X_1| = |S_1|$, essentially because $\aleph_0|S_1| = |S_1|$.

For each $b \in S_2$, choose an element $\tau(b) \in X_1$ such that $\|b-\tau(b)\| < \frac{1}{2}\sqrt{2}$. If $b \neq b' \in S_2$, then $\|b-b'\|^2 = \|b\|^2 + \|b'\|^2 = 2$, and

$$\|\tau(b)-\tau(b')\| \geq \|b-b'\| - \|b-\tau(b)\| - \|b'-\tau(b')\| > \sqrt{2} - \tfrac{1}{2}\sqrt{2} - \tfrac{1}{2}\sqrt{2} = 0,$$

so that $\tau(b) \neq \tau(b')$. Thus $\tau : S_2 \longrightarrow X_1$ is injective, and $|S_2| \leq |X_1| = |S_1|$. The converse inequality follows by symmetry. $\square$

The axiom of choice appears here both in the construction of τ and in the cardinal arithmetic. It is natural to describe the cardinal of an orthonormal basis of H as the *Hilbert dimension* of the Hilbert space H.

Definition 17.18. Let H_1 and H_2 be Hilbert spaces over $\mathbb{K}$. They are *unitarily equivalent* if there is a surjective $\mathbb{K}$-linear map $U : H_1 \longrightarrow H_2$ such that $\langle Ux, Uy\rangle = \langle x, y\rangle$ for all $x, y \in H_1$. (This condition ensures that U is also injective.)

This is the notion of "isomorphism" that is appropriate for Hilbert spaces.

Theorem 17.19. *Two Hilbert spaces H_1, H_2 are unitarily equivalent if and only if they have the same Hilbert dimension.* $\square$

§18. Linear operators in Hilbert space: generalities.

Let $T : H \longrightarrow H$ be a bounded linear operator in the Hilbert space H. Fix $y \in H$. Then, for any $x \in H$, $|\langle Tx, y\rangle| \le \|Tx\|\|y\| \le \|T\|\|x\|\|y\|$. Hence, $x \mapsto \langle Tx, y\rangle$ is a bounded linear functional on H, and so, by 13.7, there is a unique $z \in H$ such that $\langle Tx, y\rangle = \langle x, z\rangle$ for each x. Of course z depends on the given y (and on T). We write $T^*y := z$. Thus,

$$\langle Tx, y\rangle = \langle x, T^*y\rangle\,, \quad \text{and, conjugating,} \quad \langle y, Tx\rangle = \langle T^*y, x\rangle\,,$$

for any $x, y \in H$; either of these identities determines T^*.

In terms of 13.7, $Q(y) \circ T = Q(T^*y)$ for each $y \in H$. Referring to 12.10, one may express this as $T'(Q(y)) = Q(T^*y)$ or $T^* = Q^{-1} \circ T' \circ Q$. The conjugations cancel out, so that $T^* : H \longrightarrow H$ is linear.

T^* is the *adjoint* of the bounded operator T.

Definition 18.1. $T \in L(H, H)$ is *self-adjoint* or *Hermitian* if $T = T^*$. It is *skew-adjoint* or *skew-Hermitian* if $T = -T^*$.

There is a theory of adjoints for unbounded operators, in which the ideas of “self-adjoint” and “Hermitian” are distinguished from each other. But for bounded operators they coincide.

Lemma 18.2. *If* $\mathbb{K} = \mathbb{C}$, $T \in L(H, H)$ *is self-adjoint if and only if* $\langle Tx, x\rangle$ *is real for all* $x \in H$*, and is skew-adjoint if and only if* $\langle Tx, x\rangle$ *is pure imaginary for all* $x \in H$. □

Lemma 18.3. *Let* $T, T_1, T_2 \in L(H, H)$*, and* $\lambda, \mu \in \mathbb{K}$. *Then*

(a) *the identity operator and the zero operator are self-adjoint,*

(b) $(\lambda T_1 + \mu T_2)^* = \overline{\lambda} T_1^* + \overline{\mu} T_2^*$, (c) $(T_1 T_2)^* = T_2^* T_1^*$, $T^{**} = T$,

(d) T^* *is bounded (so that* $(T^*)^*$ *is defined) and* $\|T^*\| = \|T\|$,

(e) $(T^*)^* = T$, (f) $\|T^*T\| = \|T\|^2$.

Proof. *(a)–(c)* are trivial exercises. For *(d)*, suppose that $x \in H$. Then

$$\|T^*x\|^2 = \langle T^*x, T^*x\rangle = \langle TT^*x, x\rangle \le \|TT^*x\|\|x\| \le \|T\|\|T^*x\|\|x\|\,.$$

If $\|T^*x\| \ne 0$, cancel it: $\|T^*x\| \le \|T\|\|x\|$, which is also true if $\|T^*x\| = 0$. This holds for any x with $\|x\| \le 1$, so $\|T^*\| = \sup\{\|T^*x\| : \|x\| \le 1\} \le \|T\|$. Thus T^* is bounded, and *(e)* is a trivial consequence. Hence, $\|T\| = \|(T^*)^*\| \le \|T^*\|$ too. This completes the proof of *(d)*, and it follows from 10.14 that $\|T^*T\| \le \|T^*\|\|T\| = \|T\|^2$. But, for $x \in H$,

$$\|Tx\|^2 = \langle Tx, Tx\rangle = \langle T^*Tx, x\rangle \le \|T^*Tx\|\|x\| \le \|T^*T\|\|x\|^2\,,$$

so that, taking suprema over $\|x\| \le 1$, $\|T\|^2 \le \|T^*T\|$. This proves *(f)*. □

Of course one usually writes T^{**} instead of $(T^*)^*$, and there is another equality:

$$\|T\|^2 = \|T^*\|^2 = \|T^{**}T^*\| = \|TT^*\|.$$

too

Definition 18.4. The operator $T \in L(H,H)$ is *normal* if $T^*T = TT^*$.

Lemma 18.5. *When $\mathbb{K} = \mathbb{C}$, T is normal if and only if it may be expressed in the form $T = A_1 + iA_2$, where A_1, A_2 are self-adjoint bounded operators and $A_1A_2 = A_2A_1$.* □

Lemma 18.6. *When $\mathbb{K} = \mathbb{C}$, $T \in L(H,H)$ is normal if and only if $\|Tx\| = \|T^*x\|$ for all $x \in H$.* □

Proposition 18.7. *Suppose that $T \in L(H,H)$. There exists a closed subspace K of H such that $T = P_K$ if and only if T is self-adjoint and idempotent: $T = T^*$, $T^2 = T$.*

Proof. Recall from 13.6 that P_K is completely determined by

$$(\forall x \in H) \quad P_K x \in K \;\;\&\;\; x - P_K x \perp K.$$

Hence, for any $x, y \in H$, $P_K x \perp y - P_K y$ and *vice versa* and

$$\begin{aligned}\langle P_K x, y\rangle &= \langle P_K x, y - P_K y\rangle + \langle P_K x, P_K y\rangle = \langle P_K x, P_K y\rangle \\ &= \langle x - P_K x, P_K y\rangle + \langle P_K x, P_K y\rangle = \langle x, P_K y\rangle.\end{aligned}$$

This shows that $P_K^* = P_K$. As $P_K x \in K$, it follows that $P_K(P_K x) = P_K x$ for every $x \in H$, and $P_K^2 = P_K$.

(1) Suppose that $T = T^2$. Define $K := T(H)$, which is certainly a linear subspace of H. If $y = Tx \in K$, then $Ty = T(Tx) = T^2 x = Tx = y$, and, conversely, if $y = Ty$, then certainly $y \in T(H)$. Hence $K = \{y \in H : y = Ty\} = \ker(I - T)$. This shows that K is a closed subspace of H. (As $(I-T)^2 = I - 2T + T^2 = I - 2T + T = I - 2T$, the same argument shows that $L := (I - T)(H) = \ker(T)$ is a closed subspace of T.)

(2) For any $x \in H$, $x = Tx + (x - Tx)$, and, for any $k \in K$,

$$\langle x - Tx, k\rangle = \langle x, k\rangle - \langle x, T^*k\rangle = \langle x, k\rangle - \langle x, Tk\rangle = \langle x, k\rangle - \langle x, k\rangle = 0.$$

As $Tx \in K$ by definition, this shows that $T = P_K$. □

Notice that idempotence shows T is a "projection on K along L"; this idea is good in any normed space. That L and K should be orthogonal to each other is a concept only applicable in an inner product space, and is equivalent here to the self-adjointness of T.

§19. Linear operators in Banach spaces.

In a finite-dimensional vector space E, a linear operator is an isomorphism of E with itself, i.e. is invertible, either if it is surjective or if it is injective — the two conditions are equivalent. For infinite-dimensional vector spaces, this is definitely false. Elegant examples are furnished by the *shift operators* in sequence spaces. Let us take l^2, although c_0 or l^p or many other sequence spaces are equally possible. Define the *right shift* to be the mapping

$$R : l^2 \longrightarrow l^2 : (\xi_1, \xi_2, \xi_3, \dots) \mapsto (0, \xi_1, \xi_2, \dots),$$

that is, $R((\xi_n)) := (\eta_n)$, where $\eta_1 := 0$ and $\eta_{n+1} := \xi_n$ for $n \ge 1$. R is then injective, but is obviously not surjective. There is also the *left shift*:

$$L : l^2 \longrightarrow l^2 : (\xi_1, \xi_2, \xi_3, \dots) \mapsto (\xi_2, \xi_3, \xi_4, \dots),$$

which is surjective but not injective. In fact R and L are mutually adjoint.***

Lemma 19.1. *Let E and F be Banach spaces and $T : E \longrightarrow F$ a bounded linear map. Suppose T is injective. Then $T(E)$ is a closed subspace of F if and only if there is a positive number δ such that $\|Tx\| \ge \delta\|x\|$ for all $x \in E$.*

Proof. Suppose that such a δ exists. If (Tx_n) is a sequence in $T(E)$ and $Tx_n \to y$ in F, then (Tx_n) is Cauchy in F; but, for any $m, n \in \mathbb{N}$, $\|x_m - x_n\| \le \delta^{-1}\|Tx_m - Tx_n\|$, so that (x_n) is Cauchy in E. Thus $x_n \to x \in E$, and $Tx_n \to Tx$ in F, and $y = Tx \in T(E)$. This proves that $T(E)$ is closed in F.

On the other hand, if $T(E)$ is closed in F, it is a Banach space and $T : E \longrightarrow T(E)$ is a bijective and continuous. By the open mapping theorem (or, specifically, its corollary 9.9), the inverse mapping $T(E) \longrightarrow E$ is bounded. If $E = 0$, there is nothing to prove, but otherwise, let $\Delta > 0$ be the norm of this inverse mapping; then $\delta := \Delta^{-1}$ satisfies the requirements of the Lemma. □

For a further complication, consider the *weighted shift*

$$R_{[B]} : l^2 \longrightarrow l^2 : (\xi_1, \xi_2, \xi_3, \dots) \mapsto (0, \beta_1\xi_1, \beta_2\xi_2, \dots),$$

where $B := (\beta_1, \beta_2, \beta_3, \dots)$ is a bounded sequence in $\mathbb{K}$. If all the βs are non-zero, $R_{[B]}$ is injective. But, for instance, $\beta_n := 1/n$ would ensure that $\left\|R_{[B]}(e_n)\right\| \le n^{-1}\|e_n\|$. Hence, from 19.1, the image of $R_{[B]}$ is not closed. The same applies to the "multiplication operator"

$$M_{[B]} : (\xi_1, \xi_2, \xi_3, \dots) \mapsto (\beta_1\xi_1, \beta_2\xi_2, \beta_3\xi_3, \dots).$$

$M_{[B]}$ is bounded and self-adjoint, and its image is dense in l^2, since it contains all the standard basis vectors e_n, but the image is *not* closed, because of 19.1.

Definition 19.2. If E is any vector space (over any field whatever), $GL(E)$ in the algebraic context denotes the set of all bijective linear mappings $E \longrightarrow E$; that is, of all linear automorphisms of E. When E is a topological vector space (over $\mathbb{R}$ or $\mathbb{C}$), $GL(E)$ denotes the set of all continuous linear mappings $E \longrightarrow E$ which have continuous linear inverses. Such mappings are usually called "invertible".

If E is a complete linear metric space, $GL(E)$ is the set of all continuous linear bijections $E \longrightarrow E$, since, by 9.9, such a mapping has a continuous inverse.

Taking multiplication as composition of mappings, we make $GL(E)$ into a group (both in the algebraic case and in that where E has a topology).

Theorem 19.3. *Suppose that E is a Banach space.*

(a) *If $S \in L(E, E)$ and $\|S\| < 1$, then $I - S \in GL(E)$ and $\|I - S\|^{-1} \le \frac{1}{1-\|S\|}$.*

(b) *$GL(E)$ is an open set in $L(E, E)$.*

(c) *The mapping* $T \mapsto T^{-1} : GL(E) \longrightarrow GL(E)$ *is continuous.*

Proof. *(a)* Consider the series $I + S + S^2 + S^3 + \cdots$ in $L(E,E)$. Its sequence of partial sums is Cauchy, because of 10.14: if $k \le l$,

$$\left\|\sum_{n=0}^{k} S^n - \sum_{n=0}^{l} S^n\right\| = \left\|\sum_{n=k+1}^{l} S^n\right\| \le \|S^{k+1}\| + \|S^{k+2}\| + \cdots + \|S^l\|$$
$$\le \|S\|^{k+1} + \|S\|^{k+2} + \cdots + \|S\|^l < \frac{\|S\|^{k+1}}{1-\|S\|} \to 0$$

as $k, l \to \infty$. Thus, by 10.15, the series has a sum to infinity $Q := \lim_{k\to\infty} \sum_{n=0}^{k} S^n$, and

$$\|Q\| = \lim_{k\to\infty} \left\|\sum_{n=0}^{k} S^n\right\| \le \lim_{k\to\infty} \sum_{n=0}^{k} \|S\|^n = \frac{1}{1-\|S\|}. \tag{49}$$

However, also by 10.14,

$$(I-S)Q = \lim_{k\to\infty} \sum_{n=0}^{k} (I-S)S^n = \lim_{k\to\infty} (I - S^{k+1}) = I,$$

and similarly $Q(I-S) = I$. So $I - S$ is invertible with inverse Q.

(b) If $E = \{0\}$, there is nothing to prove: $GL(E) = L(E,E) = \{0\}$. Otherwise, let $T_0 \in GL(E)$. Then $T_0^{-1} \in GL(E)$, and so is a non-zero bounded linear operator. Suppose $T \in L(E,E)$ and $\|T - T_0\| < \|T_0^{-1}\|^{-1}$. Define $S := I - TT_0^{-1}$. From 10.14,

$$\|S\| = \|T_0T_0^{-1} - TT_0^{-1}\| \le \|T_0 - T\|\|T_0^{-1}\| < 1.$$

By *(a)*, $TT_0^{-1} = I - S \in GL(E)$. Hence, $T = (TT_0^{-1})T_0 \in GL(E)$.

(c) In *(a)*, $Q - I = \lim_{k\to\infty} \sum_{n=1}^{k} S^n$, and as at (49) $\|Q - I\| \le \dfrac{\|S\|}{1-\|S\|}$. In *(b)*,

$$\|T^{-1} - T_0^{-1}\| = \|T_0^{-1}\{(TT_0^{-1})^{-1} - I\}\| \le \|T_0^{-1}\|\|Q - I\|$$
$$\le \|T_0^{-1}\| \frac{\|I - TT_0^{-1}\|}{I - \|I - TT_0^{-1}\|} \le \frac{\|T_0^{-1}\|^2\|T - T_0\|}{I - \|T_0^{-1}\|\|T - T_0\|},$$

and this formula clearly shows that $T^{-1} \to T_0^{-1}$ as $T \to T_0$. □

Definition 19.4. Let E be a complex Banach space and $T \in L(E,E)$. Define the *resolvent set* $\rho(T)$ of T to be the set of those scalars $\lambda \in \mathbb{C}$ for which $\lambda I - T \in GL(E)$. Because of 0.0, $\rho(T)$ is an open subset of $\mathbb{C}$. If $\lambda \in \rho(T)$, $(\lambda I - T)^{-1}$ is defined, and the function

$$\lambda \mapsto R(\lambda; T) := (\lambda I - T)^{-1} : \rho(T) \longrightarrow L(E,E)$$

is called the *resolvent* of T. The complement of $\rho(T)$ in $\mathbb{C}$ is called the *spectrum* $\sigma(T)$ of T — that is, $\sigma(T)$ is the set of scalars λ such that $\lambda I - T$ is not invertible.

Lemma 19.5. *If* $T \in L(E,E)$, $\rho(T)$ *is open in* $\mathbb{C}$*, and* $\rho(T) \supseteq \{\lambda \in \mathbb{C} : |\lambda| > \|T\|\}$. *Furthermore,* $\|R(\lambda;T)\| \to 0$ *as* $|\lambda| \to \infty$*, and* $R(\lambda;T)$ *is continuous in* λ*.*

Proof. If $|\lambda| > \|T\|$, then $\lambda \ne 0$ and $\|\lambda^{-1}T\| < 1$. By 19.3*(a)*, $I - \lambda^{-1}T \in GL(E)$, and so $\lambda I - T = (\lambda I)(I - \lambda^{-1}T) \in GL(E)$. But also

$$\|(\lambda I - T)^{-1}\| = \|\lambda^{-1}(I - \lambda^{-1}T)^{-1}\| \le \frac{|\lambda|^{-1}}{1 - |\lambda|^{-1}\|T\|} = \frac{1}{|\lambda| - \|T\|} \to 0$$

as $|\lambda| \to \infty$.

If $\lambda I - T \in GL(E)$, by 19.3*(b)* there is some $\delta > 0$ such that $\mu I - T \in GL(E)$ whenever $|\lambda - \mu| < \delta$. This proves that $\rho(T)$ is open, and continuity of the resolvent as a function of λ follows from 19.3*(c)*. □

Corollary 19.6. *The spectrum of a bounded linear operator in a complex Banach space is a compact subset of* $\mathbb{C}$.

Proof. By the Lemma, $\sigma(T) \subseteq \{\lambda \in \mathbb{C} : |\lambda| \le \|T\|\}$ and $\mathbb{C} \setminus \sigma(T)$ is open. So $\sigma(T)$ is both closed and bounded in $\mathbb{C}$. □

Lemma 19.7. If $\lambda, \mu \in \rho(T)$, then $R(\lambda;T)R(\mu;T) = R(\mu;T)R(\lambda;T)$ and

$$R(\lambda;T) - R(\mu;T) = (\mu - \lambda)R(\mu;T)R(\lambda;T). \tag{50}$$

Proof. $(\mu I - T) - (\lambda I - T) = (\mu - \lambda)I$, and so the result follows from multiplying on the left by $R(\mu;T)$ and on the right by $R(\lambda;T)$. □

The relation (50) is called the *resolvent equation*. It leads to the conclusion that $R(\lambda;T)$ is holomorphic as a function $\rho(T) \longrightarrow L(E,E)$. To avoid discussing the deeper theory of operator-valued holomorphic functions, let us make the following

Definition 19.8. Let U be an open set in $\mathbb{C}$ and $f : U \longrightarrow L(E,E)$. f is *weakly holomorphic* on U if, for every $x \in E$ and $\phi \in E'$, the function $\lambda \mapsto \phi(f(\lambda))(x) : U \longrightarrow \mathbb{C}$ is holomorphic.

Corollary 19.9. *The mapping* $\lambda \mapsto R(\lambda;T) : \rho(T) \longrightarrow L(E,E)$ *is weakly holomorphic.*

Proof. From the resolvent equation, for $\lambda \in \rho(T)$ and sufficiently small $\zeta \in \mathbb{C}$,

$$\frac{R(\lambda + \zeta;T) - R(\lambda;T)}{\zeta} = -R(\lambda;T)R(\lambda + \zeta;T),$$

and, by 19.5, the right-hand side tends to $-R(\lambda;T)^2$ as $\zeta \to 0$. Thus, for any $x \in E$ and $\phi \in E'$, as $\zeta \to 0$

$$\frac{\phi(R(\lambda + \zeta;T)x) - \phi(R(\lambda;T)x)}{\zeta} \to \phi(-R(\lambda;T)^2 x).$$

This means that $\phi(R(\lambda;T)x)$ is holomorphic on $\rho(T)$. □

[In fact the proof shows that the resolvent is holomorphic *as an operator-valued function of a complex variable*, which at first sight is a much stronger property.]

Theorem 19.10. *The spectrum* $\sigma(T)$ *is non-empty (unless* $E = 0$*).*

Proof. Suppose that $\sigma(T) = \emptyset$. Then, for any $\phi \in E'$ and $x \in E$, $\phi(R(\lambda; T)x)$ is holomorphic $\mathbb{C} \longrightarrow \mathbb{C}$, and 19.5 shows that it tends to 0 as $|\lambda| \to \infty$. Liouville's theorem shows that it must be constantly 0. But the Hahn-Banach theorem, more specifically 12.2, shows in turn that $R(\lambda; T)x = 0$ for all λ. This is absurd, as $R(\lambda; T)$ must be invertible. $\square$

There are other proofs of the result which apply in special circumstances, for instance for self-adjoint operators or normal operators in Hilbert space.