Introduction to topological vector spaces

Bill Casselman
University of British Columbia
cass@math.ubc.ca

Infinite-dimensional spaces are ubiquitous in many branches of mathematics, and their topologies are almost always interesting. Samples of spaces that one is sure to encounter are $C_c(\mathbb{R})$, $C^\infty(\mathbb{R})$, $\mathrm{L}^p(\mathbb{R})$, and their continuous linear duals. There are two good reasons for making oneself at least somewhat familiar with the subject. On the one hand, without some preparation, it is easy to make false statements about any of these. On the other, there are many positive results in the subject that can save one an enormous amount of work.

Most of the available literature on topological vector spaces is written by enthusiasts, and I hope that a relatively short account will be valuable. My aim is here is to give an outline of techniques rather than full coverage, and from time to time explanations will be sketchy.

All vector spaces in this chapter will be complex, except perhaps in a few places where explicitly assumed otherwise.

Contents

1. Semi-norms

A Hilbert space V is a complex vector space assigned a positive definite inner product $u \bullet v$ with the property that Cauchy sequences converge. These are conceptually the simplest topological vector spaces, with the topology defined by the condition that a subset U of V is open if and if it contains a neighbourhood $\|v-u\| < \varepsilon$ for every one of the points u in U. But there are other naturally occurring spaces in which things are a bit more complicated. For example, how do you measure how close two functions in $C^\infty(\mathbb{R}/\mathbb{Z})$ are? Under what circumstances does a sequence of functions f_n in $C^\infty(\mathbb{R}/\mathbb{Z})$ converge to a function in that space? If two functions in $C^\infty(\mathbb{R}/\mathbb{Z})$ are close then their values should be close, but you should also require that their derivatives be close. So you introduce naturally an infinite number of measures of difference:

$$\|f\|_m = \sup_x \left| f^{(m)}(x) \right| ,$$

and say that $f_n \to f$ if $\|f_n - f\|_m \to 0$ for all m. The most fruitful way to put topologies on many other infinite-dimensional vector spaces is by using measures of the size of a vector that are weaker than those on Hilbert spaces.

1.1. Proposition. *If ρ is a non-negative function on the vector space V, the following are equivalent:*

(a) *for all scalars* a *and* b, $\rho(au+bv) \le |a|\rho(u) + |b|\rho(v)$;
(b) *for any scalar* a, $\rho(av) = |a|\rho(v)$, *and* $\rho(tu+(1-t)v) \le t\rho(u) + (1-t)\rho(v)$ *for all* t *in* $[0,1]$;
(c) *for any scalar* a, $\rho(av) = |a|\rho(v)$, *and* $\rho(u+v) \le \rho(u) + \rho(v)$.

The real-valued function ρ on a vector space is said to be **convex** if

$$\rho(tu+(1-t)v) \le t\rho(u) + (1-t)\rho(v)$$

for all $0 \le t \le 1$, and these conditions are essentially variations on convexity.

Any function ρ satisfying these conditions will be called a **semi-norm**. It is called a **norm** if $\rho(v) = 0$ implies $v = 0$.

A prototypical norm is the function $\|x\| = \sqrt{|x_1|^2 + \cdots + |x_n|^2}$ in $\mathbb{C}^n$, or the integral

$$\int_{\mathbb{R}^n} \left|f(x)\right|^2 dx_1 \ldots\, dx_n$$

for f in the space of continuous functions on $\mathbb{R}^n$ of compact support. The functions $\|f\|_m$ on $C^\infty(\mathbb{R}/\mathbb{Z})$ are semi-norms.

Proof. That (c) implies (b) and that (b) implies (a) is immediate. Assuming (a), we have

$$\rho(v) = \rho(a^{-1}av) \le |a|^{-1}\rho(av) \le |a|^{-1}\,|a|\rho(v) = \rho(v)\ ,$$

leading to homogeneity.

1.2. Corollary. *The kernel*

$$\ker(\rho) = \left\{v \in V \,\middle|\, \rho(v) = 0\right\}$$

of a semi-norm on a vector space is a linear subspace.

From now on I'll usually express semi-norms in norm notation—$\|v\|_\rho$ instead of $\rho(v)$.

The conditions on a semi-norm can be formulated geometrically, and in two rather different ways. The graph of a semi-norm ρ is the set Γ_ρ of pairs $(v, \|v\|_\rho)$ in $V \oplus \mathbb{R}$. Let Γ_ρ^+ be the set of pairs (v,r) with $r > \|v\|_\rho$. The conditions for ρ to be a semi-norm are that Γ_ρ^+ be convex, stable under rotations $(v,r) \mapsto (cv,r)$ for $|c| = 1$, and homogeneous with respect to multiplication by positive scalars. (I recall that a subset of a vector space is **convex** if the real line segment connecting two points in it is also in it.)

A more interesting geometric characterization of a semi-norm is in terms of the disks associated to it. If ρ is a semi-norm its open and closed disks are defined as

$$\begin{aligned} B_\rho(r-) &= \left\{v \,\middle|\, \|v\|_\rho < r\right\} \\ B_\rho(r) &= \left\{v \,\middle|\, \|v\|_\rho \le r\right\}. \end{aligned}$$

We shall see in a moment how semi-norms can be completely characterized by the subsets of V that are their unit disks. *What are the necessary conditions for a subset of a vector space to be the unit disk of a semi-norm?*

A subset of V is **balanced** if cu is in it whenever u is in it and $|c| = 1$, and **strongly balanced** if cu is in it whenever u is in it and $|c| \le 1$. (This is not standard terminology, but as often in this business there is no standard terminology.) Convex and balanced implies strongly balanced.

A subset X of V is **absorbing** if for each v in V there exists ε such that $cv \in X$ for all $|c| < \varepsilon$. It is straightforward to see that if ρ is a semi-norm then its open and closed disks $B_\rho(r-)$ and $B_\rho(r-)$ are convex, balanced, and absorbing.

A semi-norm is determined by its unit disks. If $r_v = \|v\|_\rho > 0$ then

$$\|v/r\|_\rho \begin{cases} > 1 & \text{if } r < r_v \\ = 1 & \text{if } r = r_v \\ < 1 & \text{if } r > r_v. \end{cases}$$

we have

$$\|v\|_\rho = \inf\{r > 0 \,|\, v/r \in B\}$$

for B equal to either $B_\rho(1)$ or $B_\rho(1-)$.

Conversely, suppose C to be an absorbing subset of V. The intersection of the line $\mathbb{R}\cdot v$ with V is an interval, possibly infinite, around 0. Since C is absorbing, there exists $r > 0$ such that $v/r \in C$. Define

$$\|v\|_C = \inf\{\lambda \geq 0 \,|\, v/\lambda \in C\}\,.$$

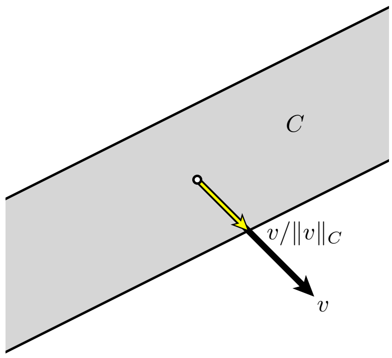

For example, if v lies on a line inside C then $\|v\|_C = 0$.

1.3. Proposition. *If C is convex, balanced, and absorbing then $\rho(v) = \|v\|_C$ is a semi-norm.*

If C is not convex or balanced then it can be replaced by its convex, balanced hull, so the first two of these requirements are not onerous. But its not being absorbing is fatal—in that case, the semi-norm will be infinite almost everywhere. So in practice it is that condition that one has to be careful about. Similarly, if one is given a formula for a semi-norm, the important thing to check is that it be finite.

Proof. It is immediate that $\|cv\|_C = c\|v\|_C$ for $c > 0$, and since C is balanced it is immediate that $\|cv\|_C = \|v\|_C$ for $|c| = 1$. Finally, the function $\rho(v) = \|v\|_C$ is convex since C is.

The semi-norm determined by C is called its **gauge**.

Several such sets C may determine the same semi-norm. For example both the open and closed unit disks of ρ determine ρ. The correspondance becomes bijective if we impose a simple condition on C. I'll call a subset of V **linearly open** if its intersection with any real line is open.

1.4. Lemma. *If ρ is a convex real-valued function then the region $\rho < c$ is linearly open.*

For example, the open unit disk defined by a semi-norm is linearly open.

Proof. Suppose P to be a point in V such that $\rho(P) < c$. I must show that every real line in V containing P contains also an open interval around P. It suffices for this to show that if Q is any other point in the vector space V, then points on the initial part of the segment from P to Q also satisfy $\rho < c$. If $\rho(Q) < c$ then the whole segment $[P, Q]$ lies in the region $\rho < c$. Otherwise say $\rho(Q) \geq c > \rho(P)$ or $\rho(Q) - \rho(P) > c - \rho(P) > 0$. We have

$$\begin{aligned} \rho\big((1-t)P + tQ\big) &\leq (1-t)\rho(P) + t\rho(Q) \\ &= \rho(P) + t\big(\rho(Q) - \rho(P)\big)\,. \end{aligned}$$

so if we choose t small enough so that

$$t\big(\rho(Q)-\rho(P)\big) < \big(c-\rho(P)\big) \quad \text{or} \quad t < \frac{c-\rho(P)}{\rho(Q)-\rho(P)}$$

then

$$\rho\big((1-t)P+tQ\big) < \rho(P) + \big(c-\rho(P)\big) = c\,.$$

1.5. Proposition. *The map associating to ρ the open unit disk $B_\rho(1-)$ where $\|v\|_\rho < 1$ is a bijection between semi-norms and subsets of V that are convex, balanced, and linearly open.*

Proof. The previous result says that if ρ is a semi-norm then $B_\rho(1-)$ is convex, balanced, and linearly open. It remains to show that if C is convex, balanced, and linearly open then it determines a semi-norm for which it is the open unit disk. For the first claim, it suffices to point out that a convex, balanced, linearly opne set is absorbing. That's because for any v the line through 0 and v must contain some open interval around 0 and inside C. This means that we can define in terms of C the semi-norm $\rho = \rho_C$.

Why is C the open unit disk for ρ? It must be shown that a point v lies in C if and only if $\|v\|_C < 1$. Since C is convex, balanced, and linearly open, the set $\{c \in \mathbb{R} \mid cv \in C\}$ is an open interval around 0 in $\mathbb{R}$. Hence if v lies in C, there exists $(1+\varepsilon)v \in C$ also, and therefore $\|v\|_C \le 1/(1+\varepsilon) < 1$.

Conversely, suppose $\|v\|_C = r < 1$. If $r = 0$, then cv lies in C for all c in $\mathbb{R}$. Otherwise, v/r is on the boundary of C—$v/\eta \in C$ for $\eta > r$ but $v/\eta \notin C$ for $\eta < r$. Since $r < 1$, there then exists some $r_+ < 1$ such that $v/r_+ \in C$, and since C is convex and v lies between 0 and v/r_+ the vector v also lies in C.

In general, linearly open sets are a very weak substitute for open ones in a vector space, but convex ones are much better behaved. Linearly open sets will occur again in the discussion of the Hahn-Banach Theorem, in which convex linearly open sets play a role. For now, I content myself with the following observation:

1.6. Proposition. *In a finite-dimensional vector space, every convex linearly open set is open.*

Proof. Let U be a linearly open subset of V. We must show that for every point of U there exists some neighbourhood contained in U. We may as well assume that point to be the origin. Let (e_i) be a basis of V. There exists $c > 0$ such that all $\pm ce_i$ are in U, and since U is their convex hull, which is a neighbourhood of the origin.

2. Topologies

A topological vector space (TVS) is a vector space assigned a topology with respect to which the vector operations are continuous. (Incidentally, the plural of "TVS" is "TVS", just as the plural of "sheep" is "sheep".) After a few preliminaries, I shall specify in addition (a) that the topology be **locally convex**, in the sense that its topology possesses a basis of neighbourhoods of 0 which are convex and (b) that the topology be **Hausdorff**.

There are two ways to define such a structure, one in terms of semi-norms and the other more directly in terms of a basis of convex neighbourhoods of 0.

There are thus two questions: *How do we define a locally convex vector space (a) by semi-norms? (b) in terms of a basis of neighbourhoods of 0?* The first method is very simple. A basis of neighbourhoods of 0 is made up of the finite intersections of open 'semi-disks' $\|v\|_\rho < \varepsilon$, and a basis of neighbourhoods of any other point is the collection of translates of those at 0. A subset is defined to be open if it contains a neighbourhood of each of its points.

We can get some idea of what neighbourhoods of 0 look like from this:

2.1. Lemma. *In a TVS, every neighbourhood of 0 (a) is absorbing and (b) contains a balanced neighbourhood of 0.*

Proof of the Lemma. Let Y be a neighbourhood of 0 in V. Given v_0 in V, $0{\cdot}v_0 = 0$, so by continuity there exists a $\delta > 0$ and a neighbourhood X of v_0 such that $cv \in Y$ if $|c| < \delta$ and $v \in X$. This shows that Y is absorbing.

Because scalar multiplication $\mathbb{C} \times V \to V$ is continuous at $(0,0)$, we can find $\varepsilon > 0$ and a neighbourhood X of 0 such that $ax \in Y$ for all $|a| < \varepsilon$ and $x \in X$. Then $\cup_{|c|\le 1} c\varepsilon X$ is still a neighbourhood of 0 contained in Y, and also balanced.

2.2. Proposition. *Given any collection of semi-norms on a vector space, vector addition and scalar multiplication are continuous in the topology defined by them. Conversely, given a locally convex topological vector space, there exists a family of semi-norms defining it.*

Proof. All is straightforward, except perhaps the definition of semi-norms. If U is a convex neighbourhood of 0, you may assume that it is balanced, since you may replace it by the union of all cU with $|c| \le 1$. This will still be a neighbourhood of 0. Now apply Proposition 1.3.

2.3. Remark. *Any complex vector space V can be given a somewhat trivial topological structure, the one with semi-norms defined by arbitrary linear functions:*

$$\|v\|_f = \big|\langle f, v\rangle\big| \,.$$

I'll call this the **linear topology**. In this topology, a basis of neighbourhoods of 0 is made up of finite intersections

$$\big|\langle f_i, v\rangle\big| < \varepsilon \text{ for } i = 1, \dots, n \,.$$

Every linear function is continuous in this topology, and every region $\mathrm{RE}(f) < c$ is open. This topology is useful in formulating results for large classes of topological vector spaces that include spaces with no interesting topological structure.

Next we look at the question, *How do we define the structure of a topological vector space by specifying a basis of neighbourhoods of 0?*

2.4. Proposition. *Suppose we are given a collection $\mathfrak{X}$ of convex, balanced, absorbing subsets of the complex vector space V satisfying these conditions:*

(a) *for every X, Y in $\mathfrak{X}$ there exists Z in $\mathfrak{X}$ with $Z \subseteq X \cap Y$;*
(b) *for every X in $\mathfrak{X}$ and $c > 0$ there exists Y in $\mathfrak{X}$ such that $Y \subseteq cX$.*

There exists on V a unique structure as topological vector space with $\mathfrak{X}$ as basis of neighbourhoods at 0. If the sets in $\mathfrak{X}$ are convex, V will be a locally convex topological vector space.

Proof. Straightforward.

In principle we know that every convex, balanced, absorbing set determines a semi-norm, but in practice it will occasionally be far simpler to work directly with $\mathfrak{X}$.

A semi-norm ρ on a TVS V is continuous if and only if the region $\rho < 1$ contains a neighbourhood of 0.

- *From now on, I shall assume unless specified otherwise that every TVS is locally convex, and I adopt the convention that a semi-norm on a locally convex TVS will be what I currently call a continuous semi-norm.*

A **basis** of continuous semi-norms is a collection of semi-norms ρ whose semi-disks $\|v\|_\rho < \varepsilon$ form a basis of neighbourhoods of 0. If F is a finite set of continuous semi-norms, then $\sup_{\rho\in F} \rho$ is also a continuous semi-norm, one whose disks are the intersections of the disks associated to the separate semi-norms in the collection. Thus if $\mathfrak{N}$ is a collection of defining semi-norms, the semi-norms $\rho_F = \sup_{\rho\in F} \rho$, as F runs through all finite subsets of $\mathfrak{N}$, form a basis of semi-norms.

Topological vector spaces offer a nice playground for those fond of logic. Theorems are usually based on intuition, but intuition alone is a poor guide to separating true from false. One cannot rely on what is valid for finite-dimensional vector spaces. Careful reasoning alone can be relied upon, but rigourous proofs are often difficult to find in an intricate maze of possibilities. Occasionally even experts allow intuition to lead them into error. One sign of skill is to find axioms for various kinds of infinite-dimensional vector spaces that are useful. One tricky question is to what extent, and when, one needs various versions of the Axiom of Choice.

Some odd features can be dealt with quickly. With no further assumption, it might happen that points in a locally convex TVS are not closed. Since $\overline{x} = x + \overline{0}$, we need only see whether or not the origin itself is closed.

2.5. Proposition. *In a locally convex TVS the closure of 0 is the intersection of the kernels of all defining semi-norms.*

In particular, the closure of 0 is a linear subspace of V.

Proof. Closed sets are the complements of open ones, so a point v lies in the complement of the closure $\overline{0}$ if and only if it has a neighbourhood that does not contain 0, or in other words if there exists a continuous semi-norm ρ such that the disk $\|u - v\|_\rho < \varepsilon$ does not contain the origin—i.e. if and only if $\|v\|_\rho > \varepsilon$ for some $\varepsilon > 0$. ▯

2.6. Corollary. *A locally convex TVS is Hausdorff if and only if $\overline{0} = 0$ or, equivalently, whenever $v \neq 0$ there exists a defining semi-norm ρ with $\|v\|_\rho \neq 0$.*

Proof. If $\|v\|_\rho > \varepsilon$ then $\|x\|_\rho = \varepsilon/2$ separates 0 from v. ▯

- *From now on, I shall assume unless specified otherwise that every TVS is Hausdorff (as well as locally convex).*

This is sometimes part of the definition of TVS, as in [Reed-Simon:1972]. The literature does not seem to be uniform.

A compact set in a TVS is one with the property that any covering of it by a collection of open sets is already covered by a finite subset of that collection. Since a single point is certainly compact, we can't expect all compact sets to be closed in a locally convex TVS unless it is Hausdorff. But that suffices:

2.7. Proposition. *Any compact set in a TVS is closed.*

Proof. It has to be shown that the complement of every compact C is open. Suppose v not in C. For each x in C we can find a continuous semi-norm ρ_x and a positive R_x such that the disks $\|y - x\|_{\rho_x} < R_x$ and $\|y - v\|_{\rho_x} < R_x$ are disjoint. Because C is compact, there exists a finite $\mathfrak{X}$ such that the disks $\|y - x\|_{\rho_x} < R_x$ for x in $\mathfrak{X}$ cover C. Let $\rho = \max_{\mathfrak{X}} \rho_x$, $R = \min_{\mathfrak{X}} R_x$. Then the disk $\|y - v\|_\rho < R$ contains no point of C. ▯

If U is a linear subspace of the vector space V, every semi-norm ρ on V determines one on U by restriction, making it into a TVS. It also induces a semi-norm on the quotient V/U:

$$\|\overline{v}\|_\rho = \inf_{v \mapsto \overline{v}} \|v\|_\rho \, .$$

Here $\inf$ means the greatest lower bound, so that (a) $\|\overline{v}\|_\rho \leq \|v\|_\rho$ for every $v \mapsto \overline{v}$ and (b) if $r > \|\overline{v}\|_\rho$ there exists $v \mapsto \overline{v}$ with $\|v\|_\rho < r$. A basis of semi-norms on a TVS defines one on the quotient.

2.8. Proposition. *A vector subspace U of the TVS V is closed if and only if the quotient topology on V/U is Hausdorff.*

Proof. The subspace U is closed if and only if each point v in the complement of U has an open neighbourhood disjoint from U, which means that there is a semi-norm ρ such that $\|u - v\|_\rho < 1$ for no point u in U. But this happens if and only if in V/U the image of this disk separates the image of v from 0. ▯

In finite dimensional topological spaces the topology can be defined in terms of convergent sequences—a set Y is closed, for example, if and only if every limit of a sequence in Y is also in Y. But in arbitrary topological spaces sequences are not sufficient to deal with convergence. This is implicit in what we have already seen—if a TVS is defined by an uncountable set of semi-norms, it doesn't seem likely that sequences can capture the structure of the neighbourhood of a point. Instead, one uses **nets**.

The motivation for the definition of a net is that a point of V is specified as the unique intersection of all of its neighbourhoods $v + X$ as X ranges over a neighbourhood basis of 0. The characteristic property of a neighbourhood basis is that if X and Y are in the basis then $X \cap Y$ contains a set in the basis—it is a **directed set**. I recall that a directed set A is a set together with an order relationship $a \preceq b$ (b is a **successor** of a), such that any two elements have a common successor. For neighbourhoods of 0, $X \preceq Y$ when $Y \subseteq X$. A net in a TVS V is a pair (A, f) where f is a map from a directed set A to V. A sequence is a special case, where $A = \mathbb{N}$ and $m \preceq n$ when $m \leq n$. A net (A, f) converges to a point v if for every neighbourhood U of v there exists an a in A such that $f(b)$ is in U for all $b \succeq a$. As suggested above, a typical example of a convergent net is that where A is a basis of neighbourhoods of a point in a locally convex TVS, and $f(a)$ for a neighbourhood a is a point of A. The reason one has to allow general nets rather than just neighbourhood bases of points is to allow one to deal with the image of a net under a continuous map.

Nets or something equivalent are necessary in dealing with convergence in arbitrary topological spaces, but there is a difficulty, in that the analogue of a subsequence is not quite intuitive. A sequence in V is a map from $\mathbb{N}$ to V, taking i to v_i. A subsequence is the composite of a monotonic, injective map n_i from $\mathbb{N}$ to $\mathbb{N}$ followed by the original sequence, all in all taking i to v_{n_i}. If (A, f) is a net in V, a **subnet** is a a triple (B, g, φ) where (i) B is a directed set; (ii) φ is a map from B to A; (iii) g is the composite of f with φ; (iv) given a in A there exists b in B with $a \preceq \varphi(b)$. To this subnet is associated the net (B, g).

With these notions, we have the standard characterization of compact sets:

2.9. Proposition. (Heine-Borel) *For a subset C of any Hausdorff topological space, the following conditions are equivalent:*

(a) if Σ is a collection of open sets covering C, then C can be covered by a finite subset of sets in Σ;
(b) any net in C possesses a convergent subnet.

3. Finite-dimensional spaces

Finite-dimensional spaces occur often in this business. It is nice to know, and should be no surprise, that they are the same familiar objects that we are used to.

3.1. Proposition. *Any Hausdorff topology on a finite-dimensional vector space with respect to which vector space operations are continuous is equivalent to the usual one.*

It is not necessary to assume the topology to be locally convex.

Proof. A basis of V determines a linear isomorphism $f: \mathbb{R}^n \to V$, which is continuous by assumption on the topology of V. By definition of continuity, if U is any neighbourhood of 0 in V there exists some disk $B(r)$ with $f\big(B(r)\big) \subseteq U$.

It remains to show that the inverse of f is continuous. For this, it suffices to show that $f\big(B(1)\big)$ contains a neighbourhood of 0.

Since the topology of V is Hausdorff, for every point s in $f\big(\mathbb{S}^{n-1}\big)$ By Lemma 2.1 we may assume U to be balanced. Since U is balanced, the neighbourhood $U/2$ is contained in U, and does not meet $s + U/2$. Since the embedding of $\mathbb{S}^{n-1}$ is continuous, we may find a neighbourhood Σ_s of s in $\mathbb{S}^{n-1}$ such that $f(\Sigma_s)$ is contained in $s + U/2$, and which does not intersect $U/2$. Since $\mathbb{S}^{n-1}$ is compact, we may find a finite number of Σ_s covering $\mathbb{S}^{n-1}$. The intersection of this finite collection of sets U_s is still a neighbourhood of 0, and does not intersect $\mathbb{S}^{n-1}$. Since it is balanced, it does not contain any points exterior to $\mathbb{S}^{n-1}$, either. So it is contained in the open disk $\|v\| < 1$. ▯

3.2. Proposition. *Any finite-dimensional subspace E of a TVS is closed.*

I am reverting here to the convention that a TVS is locally convex, although it is not a necessary assumption.

Proof. It must be shown that the complement of E in the TVS V is open, or that every point of V has a neighbourhood containing no point of E. The previous Lemma tells us that we may find a neighbourhood of 0 in V whose intersection with E is contained inside a unit sphere. There therefore exists a semi-norm ρ of V defining the topology of E. Since

$$\|e\|_\rho \le \|e-v\|_\rho + \|v\|_\rho, \quad \|e-v\|_\rho \ge \|e\|_\rho - \|v\|_\rho\,.$$

Choose $R = \|v\|_\rho$. Then for $\|e\|_\rho > 2R$

$$\|e-v\|_\rho \ge \|e\|_\rho - \|v\|_\rho > 2R - R = R\,.$$

Therefore the minimum value m of $\|e-v\|_\rho$ on the compact disk $\|e\|_\rho \le 2R$, which is at most $\|0-v\|_\rho = R$, is the minimum value on all of E. The disk $\|x-v\|_\rho < m$ is then a neighbourhood of v containing no point of E. ▯

3.3. Corollary. *If U is a closed linear subspace of a Hausdorff TVS and F a finite-dimensional subspace, then $U+F$ is closed.*

Proof. Let V be the TVS. The claim is true because the image of F in V/U is closed. ▯

3.4. Proposition. *Suppose U to be of finite codimension in the TVS V. Then U is closed if and only if every linear function on V vanishing on U is continuous.*

Proof. Let $E = V/U$. According to Proposition 2.8 and Proposition 3.1, U is closed if and only if the quotient topology on E is the usual one. But a finite-dimensional TVS is Hausdorff if and only if every linear function on it is continuous. ▯

As a special case:

3.5. Corollary. *If f is a linear function on V, it is continuous if and only if the hyperplane $f=0$ is closed.*

The following is used in the theory of partial differential equations, among other places, to verify that an eigenspace has finite dimension, and in the theory of complex analytic manifolds to verify that certain cohomology groups have finite dimension.

3.6. Proposition. *Any locally compact Hausdorff TVS is finite-dimensional.*

Proof. Let Ω be a compact neighbourhood of 0. Given any $0<c<1$ the set Ω may be covered by a finite number of $e_i + c\,\Omega$. Let E be the space spanned by the e_i, which is closed in V. Let $\overline{V} = V/E$. The image of Ω in $\overline{V}$ is contained in $c\,\Omega$, which implies that each $c^{-n}\Omega$ is contained in Ω. This implies that $\Omega = \overline{V}$. But $\{0\}$ is the only compact Hausdorff TVS. ▯

4. The Hahn-Banach theorem

- *I repeat that I include in the definition of TVS the condition that it be Hausdorff.*

A linear function f from a locally convex TVS V to $\mathbb{R}$ (or $\mathbb{C}$) is continuous if and only if there exists a semi-norm ρ of V such that $|f(v)| \le \|v\|_\rho$. Let $\widehat{V}$ be the space of all continuous linear functions on V. It is not immediately clear that $\widehat{V}$ is reasonably large—large enough to separate points in V, for example. But it is, as we'll see in this section, and the consequences of this fact are among the most basic tools in the subject. To begin with, we work with arbitrary vector spaces, not necessarily assigned a topology.

I recall that a subset of a vector space is linearly open if its intersection with any real line is open. The following Lemma is straightforward:

4.1. Lemma. *The image of a linearly open set under a surjective real linear map is linearly open.*

A **linearly open half space** (as opposed, say, to linearly closed) is a region $f > c$ where f is a real linear function. The following result about linearly open sets is the main step towards the Hahn-Banach theorem.

4.2. Proposition. (Geometric Hahn-Banach) *Suppose V to be any real vector space, $U \subseteq V$ a linear subspace and Ω a convex and linearly open subset of V whose intersection with U is contained in a linearly open half space H. There exists a linearly open half space in V containing all of Ω, whose intersection with U is H.*

Proof. By translating if necessary we may arrange that the linearly open half space in U has the origin in its boundary, or in other words is the region $f > 0$ for some linear function defined on U. So an equivalent way to state this result is this:

4.3. Proposition. (Algebraic Hahn-Banach) *If (a) U is a linear subspace of V, (b) Ω a convex, linearly open subset of V, and (c) f_U a linear function on U with $f_U > 0$ on $\Omega \cap U$, then there exists a linear function f on all of V extending f_U and > 0 on all of Ω.*

Proof. In several steps.

Step 1. Suppose that the dimension of V is two and that U is a line in V, which we may assume to be the x-axis. The intersection of Ω with the x-axis is contained in a half line starting at the origin, but does not contain the origin. Let Ω_* be the union of the multiples $\lambda\Omega$ for $\lambda > 0$. This also does not contain the origin, and is a homogeneous, open, convex cone. We may in effect replace Ω by this cone.

4.4. Lemma. *Suppose Ω to be an open, homogeneous cone in $\mathbb{R}^2$ whose intersection with some line is a half-line. It is either an open half-plane or the region lying between two half rays intersecting in an acute angle.*

Proof. We may assume the half-line to be the positive x-axis. According to Proposition 1.6, Ω is open. Let P be a point in Ω on the half-line, choose v small enough so $P_\pm = P \pm v$ are both in Ω. If the line through P_+ parallel to the x-axis does not exit Ω, then $-P_-$ will lie on that line, and Ω will contain the origin. Similarly for the parallel line through P_-. Suppose these horizontal lines from $P \pm v$ to exit Ω at points $Q_\pm$ not in Ω.

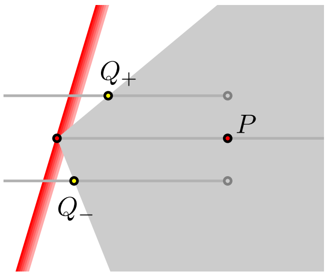

The point Q_+ will lie in the closed half plane whose boundary is the line through 0 and Q_-, again because otherwise 0 would be in Ω. If the points $Q_\pm$ are colinear, Ω must be a half plane. Otherwise Ω is the interior of the acute cone spanned by the rays through $Q_\pm$.

Step 2. We go on to the case where V is arbitrary but U has codimension one. Suppose H is the region $f_U > 0$. We can reduce to the previous case by considering the quotient of V by the kernel of f. The kernel is a subspace of U of codimension one, so this quotient is a vector space of dimension two. The image of Ω will be linearly open, and the image of U is a line, so the elementary case just done implies what we want.

Step 3. If V is finite-dimensional, the theorem follows easily by induction. Otherwise, let V and U be arbitrary vector spaces. We are looking for a linearly open half space containing Ω whose intersection with U is H. Let $\mathfrak{F}$ be the set of all pairs (U_*, H_*) with $U \subseteq U_*$, $H_* \cap U = H$, $\Omega \cap U_* \subset H_*$. This set is not empty since it contains (U, H). It can be partially ordered by double inclusion. A chain (U_α, H_α) in this order has a bounding element, namely $(\cup U_\alpha, \cup H_\alpha)$. Therefore by Zorn's Lemma there exists a maximal pair

$(U_{\max}, H_{\max})$. If $U_{\max}$ is not all of V then we can pick v not in $U_{\max}$, and by the case just done find an extension of $H_{\max}$ to the subspace spanned by $U_{\max}$ and v, which contradicts maximality.

If F is any real function on a set X, its graph Γ_F is the set of all $\big(x, F(x)\big)$. The region Γ_F^+ is the subset of $X \times \mathbb{R}$ of (x, c) where $c > F(x)$. If X is a vector space, the convexity of the set Γ_ρ^+ is equivalent to the convexity of the function ρ. Lemma 1.4 implies:

4.5. Lemma. *If ρ is convex then the region Γ_ρ^+ is linearly open.*

4.6. Corollary. *Suppose ρ to be a convex function on the real vector space V, f a linear function on a linear subspace U satisfying $f(u) \le \rho(u)$ there. Then f extends to a linear function on all of V satisfying the same inequality.*

Proof. The inequality $f \le \rho$ is equivalent to the condition $\Gamma_\rho^+ \subseteq \Gamma_f^+$.

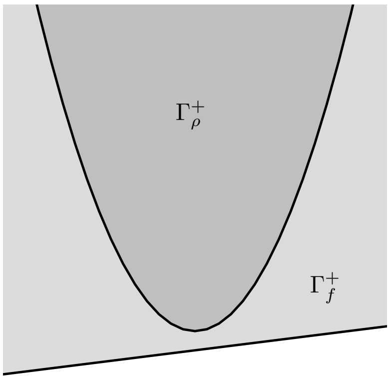

Apply the Theorem to the regions Γ_f^+ and Γ_ρ^+ inside $U \oplus \mathbb{R}$, then inside $V \oplus \mathbb{R}$.

This suffices to prove:

4.7. Theorem. (Real Hahn-Banach) *Suppose U to be a linear subspace of the real TVS V. Any continuous linear function f on U extends to one on V.*

Proof. The topology on U is induced by semi-norms on V. Since f is continuous on U, there exists a semi-norm ρ on V such that $|f(u)| \le \rho(u)$ on U. According to Corollary 4.6, there exists a linear function F extending f and satisfying the inequality $F(v) \le \rho(v)$. But $F(-v) = -F(v)$, so $-F(v) \le \rho(v)$ as well, and $|F(v)| \le \rho(v)$.

The complex case requires a simple trick.

4.8. Corollary. *Suppose ρ to be a non-negative function on the complex vector space V satisfying the inequality*

$$\rho(ax + by) \le |a|\,\rho(x) + |b|\,\rho(y)$$

whenever $|a| + |b| = 1$. (This is a complex analogue of convexity.) If f is a linear function on the subspace U satisfying $|f(u)| \le \rho(u)$ there, then f extends to a linear function on all of V satisfying the same inequality.

Proof. Note that for $|s| = 1$

$$\rho(x) = \rho(s^{-1}\,sx) \le \rho(sx) \le \rho(x)$$

which implies that $\rho(sx) = \rho(x)$. Apply the previous corollary to the real linear function $\ell(u) = \mathrm{RE}\big(f(u)\big)$ to get an extension I'll also call ℓ. Then $\ell(v) - i\ell(iv)$ is a complex linear function that agrees with f on U and extends f correctly.

4.9. Theorem. (The Hahn-Banach Theorem) *If U is a linear subspace of the locally convex TVS V, then any continuous linear function defined on U may be extended continuously to all of V.*

This is immediate from the previous Corollary.

4.10. Corollary. *If U is a closed subspace of the TVS V and $v \notin U$, there exists a continuous linear function f on V such that $f \,|\, U = 0$, $f(v) = 1$.*

Proof. According to Corollary 3.3 the vector space spanned by U and v is closed, and on this the linear function defined to be 0 on U and 1 on v is, according to Corollary 3.5, continuous. Extend it by the Hahn-Banach Theorem. ∎

4.11. Corollary. *In order for a subspace U to be dense in a TVS V, it is necessary and sufficient that every continuous linear function on V that vanishes on U vanish everywhere.*

Proof. Immediate from the preceding corollary. ∎

4.12. Corollary. *Vectors in a TVS are distinguished by values of continuous linear functionals.*

4.13. Corollary. *If U is a finite-dimensional subspace of the TVS V, there exists a continuous projection from V onto U which is the identity on U.*

Proof. The subspace U is closed by Proposition 3.2. By the Hahn-Banach theorem, coordinate functions on U may be extended to continuous linear functions on all of V. ∎

The general idea is that on a Hausdorff TVS there are lots of continuous linear functions. Here are some expansions upon this general principle:

4.14. Proposition. *Given a closed convex subset Ω of a TVS V and a point v of V not in Ω, there exists a closed affine hyperplane strictly separating v from Ω.*

Proof. We may as well assume the point to be the origin. Because Ω is closed, there exists a convex, balanced neighbourhood X not meeting it. Equivalently, the convex open set $\Omega + X$ does not contain the origin. Apply Proposition 4.3 to this convex set to obtain a real linear function f which is positive on it. If ω is a point of Ω, the open set $\omega + X$ is contained in $\Omega + X$. Since $|f(X)| < |f(v)|$, the function is continuous. ∎

4.15. Proposition. *A closed convex set Ω in a TVS is the intersection of all the closed affine half-spaces $f \le 0$ with $f \le 0$ on Ω.*

Proof. Immediate from the preceding result. ∎

The **weak topology** on a TVS V is that whose neighbourhood basis at 0 is the intersection of open cylinders $|f| < \varepsilon$ for continuous linear f.

4.16. Corollary. *A convex set Ω in the TVS V is closed if and only if it is closed in the weak topology of V.*

Proof. This follows from Proposition 4.15 once we know that a real linear function is continuous if and only if it is weakly continuous. But this is true of complex continuous functions, and the map $f \mapsto \mathrm{RE}(f)$ is a bijection between complex linear and real linear functions (see the elementary Lemma 8.6). ∎

The closed convex hull of a set X is the smallest closed convex set containing it.

4.17. Corollary. *The closed convex hull of a subset $X \subseteq V$ is the set of all x with $f(x) \le 0$ for all continuous affine functions $f \le 0$ on X.*

Similarly:

4.18. Corollary. *A closed convex cone in a TVS is the intersection of all the closed linear half-spaces $f \le 0$ with $f \le 0$ on C.*

5. Complete spaces

In a finite-dimensional TVS every Cauchy sequence converges to a unique vector—the TVS is said to be **complete**. Completeness is a crucial property in analysis, which is by tradition all about limits. The notion of completeness, as well as several related ones, can be generalized to a larger class of topological vector spaces. but there are a few technical problems.

One difficulty is that in general we must work with nets instead of sequences. This makes sense—a point in $\mathbb{R}$ is the unique intersection of all intervals of intervals $(-1/n, 1/n)$, and a point v in a vector space V is the intersection of all subsets $v + X$ as X ranges over a basis of neighbourhoods of 0 in V. A **Cauchy net** in a TVS V is a net (A, f) in V with the property that for every neighbourhood U of 0 there exists a in A such that $f(b) - f(c)$ is in U if $b, c \succeq a$. A subset U of a TVS, possibly V itself, is said to be **complete** if every Cauchy net in U converges to a point of U. If (A, f) is a Cauchy net and F is a continuous map, then $(A, F \circ f)$ is also a Cauchy net.

The next several results are special cases of results about uniform spaces (for this I refer to Chapter 6 of [Kelley:1955]).

5.1. Lemma. *A complete subset of a TVS is closed, and a closed subset of a complete TVS is complete.*

Proof. These claims follow directly from definitions. ▯

The point of completeness is this:

5.2. Theorem. *Suppose that U and V are TVS, with V complete. If A is a subset of U and f a uniformly continuous function from A to V, then f may be extended to a unique continuous function $\overline{f}$ from $\overline{A}$ to V. It, too, is uniformly continuous.*

Proof. Because the image of a Cauchy net under a continuous map is also a Cauchy net. ▯

The most common case is where A is dense in U and f is linear, in which case (as it is straightforward to verify) $\overline{f}$ is also linear.

If V is a TVS, a **completion** of V is a pair $(\iota, \overline{V})$ where $\overline{V}$ is a complete TVS and ι is an embedding of V into $\overline{V}$ such that (a) the image of V is dense in $\overline{V}$; (b) the topology of V is that induced by this embedding.

5.3. Proposition. *Every TVS possesses a completion, unique up to isomorphism.*

Proof. It will be useful later on to know what the construction is.

There is one useful type of TVS in which the complication of arbitrary Cauchy nets is not necessary. A **Cauchy sequence** in a TVS is a sequence (v_n) with the property that for every semi-norm ρ of V and $\varepsilon > 0$ there exists N such that $\|v_n - v_m\|_\rho < \varepsilon$ for $n, m > N$. A **normed TVS** is one whose topology is defined by a single norm $\|v\|$—that is to say, the sets $\|v\| < 1/n$ are a basis of neighbourhoods of 0. A **Banach space** is a normed TVS which is sequentially complete in the sense that every Cauchy sequence converges. If V is any normed TVS then its completion is no more complicated to define than to define the real numbers as completions of the rational numbers. If (u_m) and $v_m)$ are two Cauchy sequences in a normed space, they are said to be equivalent if the sequence $u_1, v_1, u_2, v_2, \ldots$ obtained from them by merging is also a Cauchy sequence, or equivalently if for every $\varepsilon > 0$ there exists N such that $\|u_m - v_m\| < \varepsilon$ for $m > N$. Define $\overline{V}$ to be the set of Cauchy sequences modulo this equivalence. I leave as an exercise to show that the natural embedding of V into $\overline{V}$ identifies it as a completion of V.

If V is an arbitrary TVS and ρ is a semi-norm, then $V_\rho = V/\ker(\rho)$ is a normed space, and $\overline{V}_\rho$ a Banach space. We can make a basis of semi-norms into a directed set by the condition $\rho \preceq \sigma$ if $\rho \le \sigma$, and then define the completion $\overline{V}$ of V to be the projective limit of the Banach spaces $\overline{V}_\rho$.

I claim that (a) this space is, up to isomorphism, independent of the choice of basis of semi-norms; (b) the embedding of V into $\overline{V}$ identifies it as a completion of V. Claim (a) is true because any two defining sets of semi-norms have a common refinement. As for (b), the first thing to verify is that V is dense in $\overline{V}$, which is an immediate consequence of the definition of a defining set of semi-norms. The second is to show $\overline{V}$ is complete. I leave this as an exercise.

If one combines Theorem 5.2 and the construction of the completion, one sees:

5.4. Proposition. *The space of continuous linear functions from a TVS to $\mathbb{C}$ and from its completion to $\mathbb{C}$ are the same.*

It is common for one TVS to be defined simply as the completion of another, and often that is the only description needed. Even where there is another description, that it is the completion of one of its subspaces is often important. For example, for many purposes all we need to know about the Hilbert space $\mathrm{L}^2(\mathbb{R})$ is that it is the completion of $C_c(\mathbb{R})$ in the norm

$$\|f\| = \left(\int_{\mathbb{R}} |f(x)|^2 \, dx \right)^{1/2} .$$

Completeness is a useful property of a TVS, but not all useful TVS are complete. For example, the continuous duals of reasonable spaces are rarely complete in natural topologies. We shall see some examples later on, in the discussion of dual spaces.

6. Metrizable and Fréchet spaces

It is often very useful to know that the topology on a TVS is that defined by a translation-invariant metric, in which case it is called **metrizable**. Metrizable TVS have many significant properties the rest do not. There is a natural criterion in terms of data we have been discussing:

6.1. Theorem. *A TVS is metrizable if and only if it possesses a countable basis of semi-norms.*

Proof. A metrizable TVS possesses as countable basis of neighbourhoods of 0 the sets where $d(0, v) < 1/n$. So the interesting point is the converse.

If V possesses a countable basis of semi-norms ρ_i then letting $\sigma_k = \sup_{i \le k} \rho_i$ we have a monotonic basis of semi-norms $\sigma_1 \le \sigma_2 \le \ldots$ Define a length on vectors:

$$|v| = \sum_1^\infty \left(\frac{1}{n^2} \right) \frac{\|v\|_{\sigma_i}}{1 + \|v\|_{\sigma_i}}$$

and then assign to V the manifestly translation invariant metric $|u - v|$. There are three things to be proven: (1) that $|u| = 0$ if and only if $u = 0$; (2) that $|u + v| \le |u| + |v|$; (3) that the topology defined by this metric is the same as that defined by the σ_i. The first is trivial. I leave the others as an exercise (see also Chapter 8 of [Trèves:1967]).

For metrizable spaces, the somewhat cumbersome machinery of nets is unnecessary to describe convergence, because sequences suffice. For example:

6.2. Corollary. *Suppose V to be metrizable. In order for a subset C of V to be closed, it is necessary and sufficient that whenever a sequence v_i in C converges in V, the limit is in C. In order for it to be compact, it is necessary and sufficient that every sequence in it possess a convergent subsequence.*

A Fréchet space is a complete TVS possessing a countable basis of neighbourhoods of 0. According to Theorem 6.1, its topology is described by a metric, and convergence by sequences. One nice feature of Fréchet spaces is heritability: closed subspaces and quotients of Fréchet by closed subspaces are again Fréchet spaces. Fréchet spaces are ubiquitous.

Example. $C^\infty(U)$

Suppose U to be an open subset of $\mathbb{R}^n$. The space $C^\infty(U)$ is that of all smooth functions on U. The semi-norms defining its topology are the

$$\|f\|_{k,\Omega} = \sup_{\Omega} \left|\partial^k f(x)/\partial x^k\right|$$

for k an n-tuple of non-negative integers, and Ω a compact subset of U. A basis is determined by an increasing sequence of sets Ω_n exhausting U.

Example. $C_c^\infty(\overline{U})$

Here U is an open relatively compact subset of $\mathbb{R}^n$. The functions in this space are the smooth functions on all of $\mathbb{R}^n$ with support in $\overline{U}$. Semi-norms:

$$\|f\|_k = \sup_U \left|\partial^k f(x)/\partial x^k\right|$$

for k an n-tuple of non-negative integers.

Example. $C^\infty(M)$

Here M is a compact n-dimensional manifold. Each point of M has a neighbourhood that looks like a relatively compact open subset of $\mathbb{R}^n$, and there is a unique locally convex topology on $C^\infty(M)$ agreeing with the topology defined in the previous paragraph on the corresponding images of the spaces $C^\infty(\overline{U})$. This can be seen easily by using partitions of unity on M.

Example. $\mathcal{S}(\mathbb{R}^n)$

This is the **Schwartz space** of $\mathbb{R}^n$, that of all smooth functions f on $\mathbb{R}^n$ such that

$$\|f\|_{k,m} = \sup_{\mathbb{R}^n} \big(1+|x|\big)^k \partial^m f(x)/\partial x^m$$

is finite, where k and m range over n-tuples of natural numbers non-zero integers.

There is an interesting relationship between the last two examples. There is an analogue of stereographic projection, mapping $\mathbb{R}^n$ into the n-sphere $\mathbb{S}^n$ in $\mathbb{R}^{n+1}$ and identifying it with one of the poles. The recipe is that a point x in $\mathbb{R}^n$ is taken to the point $(sx, 1-s)$ on the sphere $|sx|^2 + z^2 = 1$ intersected by the line from $(x, 0)$ to $(0, 1)$:

$$x \longmapsto \left(\frac{2x}{1+|x|^2}, \frac{1-|x|^2}{1+|x|^2}\right) .$$

Any smooth function on $\mathbb{S}^n$ thus pulls back to a function on $\mathbb{R}^n$.

6.3. Proposition. *The pull back via stereographic projection of a smooth function on $\mathbb{S}^n$ lies in the Schwartz space $\mathcal{S}(\mathbb{R}^n)$ if and only if it vanishes of infinite order at the north pole* $(0, 1)$.

The first step in proving this is to note that the functions $t_i = x_i / \sum x_i^2$ are in effect parameters at the pole on $\mathbb{S}^n$. The theorem can then be verified by showing that

$$\partial^m/\partial t^m = \sum P_{m,k}(x)\, \partial^k/\partial x^k$$

where k and m are multi-indices.

The idea of analyzing growth conditions at ∞ by making them local on some extended space is a valuable one.

Suppose U and V to be TVS, F a linear map $F\colon U \to V$. The function F is continuous if it satisfies this condition: *for every semi-norm σ on V and $\varepsilon > 0$ there exists a semi-norm ρ on U and $\delta > 0$ such that $\|F(u)\|_\sigma < \varepsilon$ whenever $\|u\|_\rho < \delta$.* Scaling the semi-norms, this is equivalent to the condition that *for every semi-norm σ on U there exists a semi-norm ρ on U such that $\|u\|_\rho < 1$ implies that $\|F(u)\|_\sigma < 1$*, or yet again equivalently *every semi-norm σ on V there exists a semi-norm ρ on U such that $\|F(u)\|_\sigma \le \|u\|_\rho$.* For finite-dimensional spaces, every linear map is continuous, but this is far from true for infinite-dimensional ones. In fact, even being continuous is a relatively weak condition—you might also want that it be a **homomorphism**, which means that the quotient topology on $F(U)$ is the same as that induced from the embedding of $F(U) \subseteq V$; or that it have **closed range**, which means that $F(U)$ is a closed linear subspace of V. Fréchet spaces have one useful property that makes them look pleasantly like finite-dimensional spaces. For Fréchet spaces, these last two conditions are the same:

6.4. Proposition. (Open mapping theorem) *A continuous linear map $T\colon U \to V$ of Fréchet spaces has closed range if and only if it is a homomorphism.*

Proof. Suppose that T has closed range. It may be assumed to be surjective. It has to be shown that the image under T of any disk $B = B_{\rho,r}$ in U (where ρ is a semi-norm) contains a neighbourhood of 0 in the image space. In order to prove this, we have to use the Baire Category Theorem, which asserts that if a countable union of closed subsets of a complete metric space is never the union of a countable union of nowhere dense sets. In our situation, the union of the sets $T(nB) = nT(B)$ is all of V, so the closure of one of them, say $nT(B)$, contains some open set Ω. But then $\Omega - \Omega$ will also be open, contain 0, and be in the closure of the image of $2nB$. Scaling, this tells us that the closure of $T(B)$ contains a neighbourhood of 0.

Now we switch to using the topologies in both spaces defined by a translation-invariant metric. For each r let B_r be the metric disk $|v| < r$. We know at this point that for every $r > 0$ there exists $\rho > 0$ with the closure of $T(B_r)$ containing B^*_ρ. Given one choice ofρ, a smaller choice will be equally valid.

We must show that each $T(B_r)$ contains some B_ρ. Fix r, choose some $0 < q < 1$, and let $R = r(1-q)$, or

$$r = R(1 + q + q^2 + \cdots)$$

For each $k > 1$ let ρ_k be such that B_{ρ_k} is contained in $\overline{T(B_{Rq^k})}$. We may assume that $\rho_{k+1} < \rho_k$ and that $\rho_k \to 0$ since we cam always make them smaller.

I claim that $T(B_r) \supset B_{\rho_0}$. Suppose y to be in B_{ρ_0}.

Let $y_0 = y$. We know that $\overline{T(B_{R_0})}$ is dense in B_{ρ_0}, so we can find x_0 in B_{R_0} such that $|T(x_0) - y| < \rho_1$. Let $y_1 = y - T(x_0)$. Since y_1 lies in B_{ρ_1}, we can find x_1 in B_{rq} such that $|y_1 - T(x_1| < \rho_2$. By induction we may choose x_k in B_{rq^k} such that $y_{k+1} = y_k - T(x_k)$ lies in $B_{\rho_{k+1}}$. But then

$$y = y_k + T(x_0) + T(x_1) + \cdots + T_{x_k}, \quad |x_k| < rq^k, \quad |y_k| < \rho_{k+1}\,.$$

Therefore the series

$$x = x_0 + x_1 + x_2 + \cdots$$

converges in U to an element of B_r and $T(x) = y$.

Now suppose that T is a homomorphism. The quotient topology on $T(U)$ makes it a Fréchet space, hence complete, and since its topology is also that induced from its embedding into V it is complete in the embedding topology. Hence it is closed in V. ▯

An exact sequence of TVS is an exact sequence in the algebraic sense

$$0 \to U \to V \to W \to 0\,,$$

in which the maps are linear and continuous. I call it a **topological exact sequence** if quotients are isomorphic to images—i.e. if the maps are homomorphisms.

6.5. Corollary. *Any exact sequence of Fréchet spaces is topologically exact.*

This is a first step towards finding a large and useful category of TVS that behave almost as well as finite-dimensional vector spaces. Another way to phrase results of this section is to say that the category of Fréchet spaces together with homomorphisms of TVS as its morphisms form an abelian category.

I should add that although these results are very pretty, it is often quite difficult to verify that a map has closed range (or, equivalently, that it is a homomorphism). Often such a theorem is in fact a major accomplishment. However, there are some useful examples that are not too complicated, variants of a well known theorem of Émile Borel.

7. Weakly bounded means strongly bounded

Corollary 4.16 tells us that a convex set that is closed in the weak topology is also closed in the original topology. The following result is much deeper, but perhaps not a complete surprise:

7.1. Theorem. *A subset of the TVS V is bounded if and only if for each continuous linear function f on V the function $|\langle f, v\rangle|$ is bounded on V.*

In other words, the subset is bounded if and only if it is weakly bounded.

This is V.23 of [Reed-Simon:1972]. This entire section will be devoted to proving it, and since their exposition is quite clear, I refer to them for details of the proof. The hard part is proving the claim for Banach spaces.

BANACH-STEINHAUS. If V is a Banach space, a continuous linear function on V is a linear function such that

$$|f(x)| \le C\|x\|$$

for some $C > 0$, so we may define

$$\|f\| = \sup_{\|x\|=1} \frac{|f(x)|}{\|x\|} .$$

This is a norm on $\widehat{V}$.

7.2. Lemma. *If V is a Banach space, so is $\widehat{V}$.*

Proof. It must be shown that $\widehat{V}$ is complete. This means that a Cauchy sequence of functions f_n in $\widehat{V}$ converge in norm to a bounded linear function f. The limit must be defined, shown to be bounded, and shown to be the norm limit of the f_n. I leave these claims as exercises. (See [Reed-Simon:1972], III.2.)

Any v in V determines a continuous linear function on $\widehat{V}$, giving in effect a linear map from V to the continuous linear dual of $\widehat{V}$. It rarely haopens that this is an isomorphism, but we do have:

7.3. Lemma. *If V is a Banach space, the canonical map from V to the continuous linear dual of $\widehat{V}$ is an isometric embedding.*

Proof. I leave this as an exercise, too. (See [Reed-Simon:1972], III.4.)

7.4. Proposition. (Banach-Steinhaus) *Suppose V a Banach space, $\mathcal{F}$ a set of continuous linear functions in $\widehat{V}$. Suppose that for each v in V the set of all $|\langle F, v\rangle|$ for F in $\mathcal{F}$ is bounded. Then the set $\mathcal{F}$ is bounded.*

Proof. Also an exercise. The standard proof uses the Baire Category Theorem. (See [Reed-Simon:1972], III.8–9.)

7.5. Corollary. *Suppose V to be a Banach space, Ω a subset of V satisfying the condition that for every $\widehat{v}$ in $\widehat{V}$ the set $|\langle \widehat{v}, v\rangle|$ as v ranges over $\mathcal{V}$ is bounded. Then Ω is bounded.*

Proof. Apply the combination of Lemma 7.3 and Proposition 7.4 with $\widehat{E}$ replacing V in the latter.

PROOF OF THE THEOREM. The proof of Theorem 7.1 is now straightforward. Let Ω be a weakly bounded subset of V. It must be shown that for every semi-norm ρ the values of $\|v\|_\rho$ are bounded on Ω. Let V_ρ be the quotient of V by the kernel of ρ, and $E = \overline{V}_\rho$ be its completion, which is a Banach space. Apply the previous result.

8. Duals

For every TVS V I define $\widehat{V}$ to be the vector space of continuous linear maps from V to $\mathbb{C}$. (This is, once again, not standard notation, but as elsewhere there is no single accepted convention. Many Americans and all French write V' but others, including perhaps all Russians, go for V^*. Me, I am neither French nor Russian, and I use V^* for the conjugate linear dual.)

The dual is well defined by the space V and its topology, but several natural questions arise here that do not arise for finite-dimensional spaces:

- *What topology is to be assigned to $\widehat{V}$?*
- *To what extent does $\widehat{V}$ determine the space V?*
- *What properties of V can be deduced from those of $\widehat{V}$?*

Duals are often easier to work with than the original TVS. For example, there exists a fundamental criterion, in terms of dual spaces, for a continuous map from one Fréchet space to another to be a homomorphism. It is a relatively straightforward application of this criterion to prove Émile Borel's theorem that every power series in x is the Taylor series of some smooth function on $\mathbb{R}$. The proof uses the dual of the vector space of formal power series, which is isomorphic to the vector space of polynomials in x—for most people a more comfortable object to deal with.

Also, in many cases, important spaces are defined—or at least characterized—in terms of duality. This includes, in a straightforward way, distributions, but also includes the π- and ε-tensor products.

DUAL PAIRS. It is best to treat V and $\widehat{V}$ as symmetrically as possible, and to do this I look at a slightly more abstract situation. Suppose V, $\mathring{V}$ to be a pair of complex vector spaces together with a bilinear map $(v, \mathring{v}) \longmapsto \langle \mathring{v}, v\rangle$ from $V \times \mathring{V}$ to $\mathbb{C}$. Any fixed $\mathring{v}$ in $\mathring{V}$ determines a linear function

$$\varphi_{\mathring{v}} \colon V \to \mathbb{C}, \quad v \longmapsto \langle \mathring{v}, v\rangle$$

on V, and similarly each v determines a linear function φ_v on $\mathring{V}$. These give linear maps from $\mathring{V}$ to

$$\widetilde{V} = \text{ the algebraic linear dual of } V ,$$

and also from V to the algebraic dual of $\mathring{V}$. I say that the pair together with the bilinear pairing form a **dual pair** if each of these maps is an embedding, or more informally if each of V, $\mathring{V}$ separates points in the other.

From now on, unless specified otherwise, V and $\mathring{V}$ will always be a dual pair.

Functions in $\mathring{V}$ determine on V the structure of TVS endowed with a locally convex Hausdorff topology $\sigma(V, \mathring{V})$. It is called the **weak topology** determined by $\mathring{V}$, and it has as basis of semi-norms

$$\|v\|_F = \sup_{\mathring{v} \in F} |\langle \mathring{v}, v\rangle|$$

where F is a finite subset of $\mathring{V}$. Symmetrically, those in V define the weak topology $\sigma(\mathring{V}, V)$ on $\mathring{V}$.

8.1. Proposition. *If $\mathring{V}$ is assigned the weak topology $\sigma(\mathring{V}, V)$, the embedding of V into the algebraic dual of $\mathring{V}$ identifies it with the continuous linear dual of $\mathring{V}$.*

That is to say, every linear function on $\mathring{V}$ that is weakly continuous lies in V. For reasons of symmetry, the same holds for $\mathring{V}$ and V.

Proof. A linear function f is continuous in the topology $\sigma(V, \mathring{V})$ if and only if $|f| < 1$ on some region where $|\langle \mathring{v}_i, v\rangle| < 1$ for a finite number of $\mathring{v}_i$ in $\mathring{V}$. Let U be the linear subspace of V where all $\mathring{v}_i = 0$. If f is not 0 on U then it is unbounded there, which is a contradiction. So f factors through the quotient V/U, and f must be a linear combination of these $\mathring{v}_i$, hence in $\mathring{V}$. ∎

The Hahn-Banach theorem and the definition of $\widehat{V}$ imply that if V is endowed with an initial structure of TVS, then $(V, \widehat{V})$ make up a dual pair. The initial topology of V is finer than the corresponding weak topology in the sense that weakly open sets in V are open in the initial topology. A closed subset in V is clearly weakly closed with respect to $\mathring{V}$, and according to Theorem 7.1 a subset of V is bounded if and only if it is weakly bounded. These considerations apply in particular to an arbitrary complex vector space and its algebraic dual, the first endowed with what I have called elsewhere the linear topology.

The space $\widehat{V}$, together with its weak topology, is traditionally expressed as $\widehat{V}_\sigma$, but I'll express it as $\widehat{V}_\omega$. There exists a canonical map from V to the continuous linear dual of $\widehat{V}_\omega$, taking v to $\widehat{\widehat{v}}$, where

$$\langle \widehat{\widehat{v}}, \widehat{v}\rangle = \langle \widehat{v}, v\rangle \, .$$

In a somewhat different language, the previous Proposition asserts that:

8.2. Proposition. *The map $v \mapsto \widehat{\widehat{v}}$ from V to the continuous dual of $\widehat{V}_\sigma$ is a linear isomorphism.*

If $F\colon U \to V$ is a linear map and $f\colon V \to \mathbb{C}$ is a linear map, then $f\circ F$ is also linear. I say that, by definition, $F\colon U \to V$ is weakly continuous (with respect to the given pairings) if $f\circ F$ lies in $\mathring{U}$ for every f in $\mathring{V}$. This composition is called the adjoint $\mathring{F}$ of F, and another way to formulate its definition is by the equation

$$\langle \mathring{F}(\mathring{v}), u\rangle = \langle \mathring{v}, F(u)\rangle \, .$$

Let me make this more reasonable, by pointing out that if U and V is a TVS and $\mathring{U} = \widehat{U}$, $\mathring{V} = \widehat{V}$, every continuous map is weakly continuous. This is trivial to verify.

POLARITY. Assume V, $\mathring{V}$ to be a dual pair. If X is any subset of V, its **polar** in $\mathring{V}$ is the subset

$$X^\circ = \left\{ \mathring{v} \in \mathring{V} \,\middle|\, |\langle \mathring{v}, x\rangle| \le 1 \text{ for all } x \in X \right\} .$$

(In older literature, this is called the **absolute polar**.) For example, for Euclidean space paired with itself the polar of the ball of radius r is the closed ball of radius $1/r$. The polar of the band $|x_n| < 1$ in Euclidean space of n dimensions is the line segment $|x_n| \le 1$ in the x_n-axis.

There are a few basic properties of this construction worth stating, all of them easy to verify.

8.3. Proposition. *Polars have the following properties:*

(a) every polar X° is closed, balanced, and convex;
(b) a set and its closed, balanced, convex hull have the same polar;
(c) $X \subseteq Y$ if and only if $Y^\circ \subseteq X^\circ$;
(d) $(X \cup Y)^\circ = X^\circ \cap Y^\circ$;
(e) if X is bounded in the weak topology then X° is absorbing.

8.4. Proposition. *If U is a linear subspace of V then*

$$U^\circ = \left\{ \mathring{v} \in \mathring{V} \,\middle|\, \langle \mathring{v}, u\rangle = 0 \text{ for all } u \in U \right\} .$$

This is the **annihilator** of U in $\overset{\circ}{V}$.

Proof. Suppose u in U and $\mathring{v}$ in U°. Then $|\langle \mathring{v}, cu\rangle| = c|\langle \mathring{v}, u\rangle| < 1$ for all $c > 0$, hence $\langle \mathring{v}, u\rangle = 0$.

The polar can be interpreted in terms of real convex hulls. First a few very elementary lemmas.

8.5. Lemma. *The region $|z| \leq 1$ in $\mathbb{C}$ is the intersection of the regions $\mathrm{RE}(\zeta z) \leq 1$ as ζ ranges over the unit circle in $\mathbb{C}$.*

8.6. Lemma. *The map $f \mapsto \mathrm{RE}(f)$ is a bijection between complex linear and real linear functions on V.*

Proof. If $f(v) = a(v) + ib(v)$ is a complex linear function then $f(iv) = if(v) = -b(v) + ia(v)$, so $b(v) = -a(iv)$. Thus we can reconstruct f from a, and the map from f to its real part is injective. Conversely, given a real linear function a we set $f(v) = a(v) - ia(iv)$.

Combining these two:

8.7. Lemma. *The polar of X is the set of all complex linear functions f in $\overset{\circ}{V}$ with the property that $\mathrm{RE}(f) \leq 1$ on the closed, convex, balanced hull of X.*

8.8. Proposition. *For X in V, the weakly closed, convex, balanced hull of X is the intersection of all regions $|\langle \mathring{x}, x\rangle| \leq 1$ as $\mathring{x}$ ranges over X°.*

Proof. By Proposition 4.15, the weakly closed, convex, balanced hull of X is the intersection of all affine half-planes $f \leq 0$ for real continuous affine functions f such that $f(\zeta x) \leq 0$ on X. But since X is a neighbourhood of 0, no linear function can be ≤ 0 on X, so every such f can be scaled to be of the form $\mathrm{RE}(\zeta\widehat{v}) - 1$ for $\widehat{v}$ in X°. Thus X is the same as the region where $|\langle \widehat{v}, x\rangle| \leq 1$ for all $\widehat{v}$ in X°.

By Proposition 8.1, V may be identified with the weak dual of $\overset{\circ}{V}$. The following is therefore just a restatement of the Proposition.

8.9. Corollary. *For $X \subseteq V$, its double polar $X^{\circ\circ}$ is the weakly closed, balanced, convex hull of X.*

From the Proposition and Proposition 8.4:

8.10. Corollary. *If U is a linear subspace of V, $U^{\circ\circ}$ is its closure.*

In summary:

8.11. Proposition. *The map $X \mapsto X^\circ$ is a bijection of the weakly closed, balanced, convex subsets of V with those of $\overset{\circ}{V}$.*

Recall that if ρ is a semi-norm on V then $B_\rho(1)$ is the set of all v with $\|v\|_\rho \leq 1$. The following is a straightforward exercise.

8.12. Lemma. *If ρ is a semi-norm on V and $X = B_\rho(1)$ then*

$$\|v\|_\rho = \sup_{\widehat{v} \in X^\circ} |\langle \widehat{v}, v\rangle| \,.$$

The polars of neighbourhoods of 0 in V are of particular interest.

8.13. Lemma. *If V is a TVS, then $\widehat{V}$ is the union in $\widetilde{V}$ of polars X°, as X ranges over neighbourhoods of 0 in V.*

Proof. For X a neighbourhood of 0 in V, its polar in $\widetilde{V}$ is the set of all linear functions $\widehat{v}$ on V such that $|\langle \widehat{v}, x\rangle| \leq 1$ for all x in X. But by definition every $\widehat{v}$ lies in such a set.

TOPOLOGIES ON THE DUAL. Now suppose V to be a TVS and $\overset{\circ}{V} = \widehat{V}$. So far, we have seen only the weak topology $\sigma(\widehat{V}, V)$ on $\widehat{V}$. But there are several other useful ones, determined by certain families $\mathfrak{S}$ of bounded sets in V.

A **covering** of V is any collection of V that . . . well . . . covers V—i.e. the union of all the sets in the collection is V itself. I call a covering $\mathfrak{S}$ **adequate** if it satisfies these three conditions:

($\mathfrak{S}_0$) the sets in $\mathfrak{S}$ are bounded;
($\mathfrak{S}_I$) if X, Y are in $\mathfrak{S}$ then there exists Z in $\mathfrak{S}$ with $X \cup Y \subset Z$;
($\mathfrak{S}_{II}$) if $\lambda \in \mathbb{C}$ and X in $\mathfrak{S}$ there exists Y in $\mathfrak{S}$ such that $\lambda X \subset Y$.

The topology determined by $\mathfrak{S}$ on $\widehat{V}$ is defined by specifying as basis of semi-norms the functions

$$\|\widehat{v}\|_X = \sup_{x \in X} \left|\langle \widehat{v}, x\rangle\right|$$

for X in $\mathfrak{S}$. This is well defined since continuous linear functions are bounded on bounded sets. A set X and its closed, balanced, convex hull determine the same semi-norm.

Frequently used examples of $\mathfrak{S}$ are the finite subsets of V, which give the weak topology on V, and the bounded subsets of V, which give the **strong topology**.

8.14. Proposition. *If $\mathfrak{S}$ is adequate, the disks*

$$B_{X,\varepsilon} = \left\{\widehat{v} \in \widehat{V} \;\middle|\; \|\widehat{v}\|_X < \varepsilon\right\} \quad (X \in \mathfrak{S})$$

form a basis of neighbourhoods of 0 in the $\mathfrak{S}$-topology.

Proof. The conditions laid out in Proposition 2.4 correspond exactly to those for $\mathfrak{S}$.

This topology is called the topology of **uniform convergence** on sets in $\mathfrak{S}$. Two collections $\mathfrak{S}$ and $\mathfrak{T}$ determine the same topology if and only if they are cofinal with respect to each other. The space $\widehat{V}$ with this topology is expressed as $\widehat{V}_{\mathfrak{S}}$. I'll repeat for emphasis: *the space of continuous linear functions from V to $\mathbb{C}$ is $\widehat{V}$, and this space together with the particular choice of topology determined by $\mathfrak{S}$ is $\widehat{V}_{\mathfrak{S}}$.* The weak topology, as I recall, is that corresponding to the set of finite subsets of V and denoted as $\widehat{V}_{\omega}$.

According to Lemma 8.13, the set of all X° for X a neighbourhood of 0 in V is an adequate cover of $\widehat{V}$.

The following is just a restatement of Lemma 8.12:

8.15. Theorem. *The topology on a TVS V is that determined by the set $\mathfrak{S}$ of all polars $X^\circ \subset \widehat{V}$ as X ranges over all neighbourhoods of 0 in V.*

This is a step towards the idea that the notion of $\mathfrak{S}$-topologies is sufficient—that there might not be any reason to look for others.

EQUICONTINUOUS SETS. A linear function $\widehat{v}$ on V is continuous if there exists a neighbourhood X of 0 in V such that $|\langle \widehat{v}, x\rangle| \le 1$ for every x in X. A subset E of $\widehat{V}$ is called **equicontinuous** if there exists a neighbourhood X of 0 in V such that $|\langle \widehat{v}, x\rangle| \le 1$ for all $\widehat{v}$ in E, x in X or, equivalently, if it is contained in the polar X° of a neighbourhood X of 0 in V.

Motivation for the name:

8.16. Lemma. *If E is an equicontinuous subset of $\widehat{V}$ and B a bounded subset of V, then for some $C > 0$ and all $\widehat{v}$ in E, b in B we have $\left|\langle \widehat{v}, b\rangle\right| \le C$.*

In other words, an equicontinuous subset of $\widehat{V}$ is bounded for the strong topology.

Proof. Suppose $E \subseteq X^\circ$. Since B is bounded, $B \subset cX$ for some $c > 0$. But then $|\langle \widehat{v}, b\rangle| \le c$ for all $\widehat{v}$ in E, b in B.

8.17. Proposition. *The topology of a TVS V is that determined by the set $\mathfrak{S}$ of all equicontinuous sets in $\widehat{V}$.*

Proof. Lemma 8.12 says that the topology on V is determined by the set of polars, and the definition of an equicontinuous set is that it be a subset of a polar of a neighbourhood of 0. The $\mathfrak{S}$-topology is not changed by adding to $\mathfrak{S}$ the subsets of sets in $\mathfrak{S}$.

The strong topology on $\widehat{V}$ is defined by the collection of all bounded subsets of V. I denote the space together with this topology as $\widehat{V}_\beta$.

There exists a canonical injection of the algebraic dual $\widetilde{V}$ into the topological direct product $\mathbb{C}^V$, the set of all functions from V to $\mathbb{C}$, taking $\widetilde{v}$ to $\big(\langle \widetilde{v}, v\rangle\big)$. The direct product has a canonical topology, whose open sets are the products (X_v) with X_v open in $\mathbb{C}$ for all v and $X_v = \mathbb{C}$ for all but a finite number of v. The restriction to $\widehat{V}$ of the injection of $\widetilde{V}$ into $\mathbb{C}^V$ is continuous in the weak topology, by definition of that topology. The image of $\widetilde{V}$ consists of all (x_v) satisfying conditions of linearity (a) $x_{u+v} = x_u + x_v$ and (b) $x_{cv} = cx_v$. It is closed, because each linearity condition is on only a finite number of coordinates. In other words, the space of linear maps from V to $\mathbb{C}$ may be identified with a closed subspace of $\mathbb{C}^V$.

So we have inclusions

$$\widehat{V} \subseteq \widetilde{V} \subset \mathbb{C}^V ,$$

The image of $\widehat{V}$ in $\mathbb{C}^V$ is not necessarily closed. The closure of a subset of $\widehat{V}$ may contain points of $\mathbb{C}^V$ that do not come from continuous maps. However, according to Lemma 8.13 the image of $\widehat{V}$ in $\widetilde{V}$ is the union of polars X° for X a neighbourhood of 0 in V, and:

8.18. Lemma. *If X is a neighbourhood off 0 in V then the closure of the image of X° is also contained in the image of X°.*

In particular, the closure of every equicontinuous subset off $\widehat{V}$ is contained in the image off $\widehat{V}$.

Proof. Let X be a neighbourhood of 0 in V, $E = X^\circ$. For each c in $\overline{E}$ we want to show that $|c_x| \le 1$ for all x in X. Since E is equicontinuous we may find a neighbourhood X such that $|\langle \widehat{v}, x\rangle| < 1/2$ for all $\widehat{v}$ in E. Since c is in the closure of E we may find for each x in X some $\widehat{v}$ in E such that $|c_x - \langle \widehat{v}, x\rangle| < 1/2$. But then

$$|c_x| < |\langle \widehat{v}, x\rangle| + 1/2 < 1 .$$

Any finite set of points in $\widehat{V}_\sigma$ is compact, but *a priori* it is not evident that there are any more interesting ones. The following result guarantees that there are lots of them.

8.19. Proposition. (The Banach-Alaoglu theorem) *The weak closure of an equicontinuous subset of $\widehat{V}$ is weakly compact.*

In particular, if X is a neighbourhood of 0 in V then X° is compact in $\widehat{V}$. As a consequence of this result and Proposition 8.17, the topology on V is determined by compact subsets of $\widehat{V}$.

Proof. Let E be an equicontinuous subset of $\widehat{V}$. Its closure in $\widehat{V}_\sigma$ may be identified with its closure in $\mathbb{C}^V$. For each v in V let $E(v)$ be the image of the set of values $\langle \widehat{v}, v\rangle$ in $\mathbb{C}_v$, as $\widehat{v}$ ranges over the closure of E. By Lemma 8.16, the set $E(v)$ is compact. A well known theorem of Tychonoff implies that the product of compact sets is compact, so the closure of E is compact.

9. Distributions

One important space is $C_c^\infty(U)$ where now U is any open subset of $\mathbb{R}^n$. It is the union of all the spaces $C_c^\infty(\overline{\Omega})$, where Ω is a relatively compact open set with $\overline{\Omega} \subseteq U$. It therefore fits into the scheme now to be described.

An **LF-sequence** is a sequence of Fréchet spaces

$$V_1 \subseteq V_2 \subseteq V_3 \subseteq \ldots$$

where each one is embedded as a closed subspace of the next. A locally convex topology is assigned to the union V_* of the V_i according to the following specification: *A convex subset U of the union is defined to be a neighbourhood of 0 if and only if $U \cap V_i$ is a neighbourhood of 0 in V_i.* This definition is unusual in these notes, since the topology is not defined by semi-norms. In principle it would be possible to do so, but would not be a natural step to take. It has to be verified that this specification of the topology makes the union into a topological vector space, but that's easy. What's not so easy is seeing that the V_i are closed subspaces of V_*, and that the induced topology is the original one. This is assured by:

9.1. Proposition. *Suppose V to be a closed vector subspace of the TVS V_*, X to be a convex neighbourhood of 0 in V. For any x_* in V_* not in X, there exists a convex neighbourhood X_* of 0 in V_* such that $V \cap X_* = U$, and not containing x_*.*

Proof. Because V is closed in V_* we can find Ω_* be a neighbourhood of 0 in V_* such that $V \cap \Omega_* \subseteq X$. The convex hull X_{**} of Ω_* and X will be a neighbourhood of 0 in V_* with $X_{**} \cap V = X$. It might contain v_*, though. Consider the quotient V_*/V, a TVS. If the image of v_* is 0 in it, then v_* lies in V, hence not in X_*. Otherwise, we can find a neighbourhood Y of 0 in V_*/V not containing the image of v_*. Let $X_* = f^{-1}(Y) \cap X_{**}$. ▯

The standard example is $C_c^\infty(U)$ with U open in $\mathbb{R}^n$. In this case, for each i let Ω_k be the subset of U of points at distance greater than $1/2^k$ from the complement of U. Set $V_k = C_c^\infty(\overline{\Omega}_k)$. This satisfies all the criteria of Proposition 9.1. We'll explore this example in more detail a bit later on.

The most important result about LF-spaces is this, which follows easily from the previous result:

9.2. Proposition. *If W is a TVS then a linear map from V_* to W is continuous if and only if it is continuous on each V_k.*

9.3. Proposition. *The space V_* is complete.*

If $V = C_c^\infty(\mathbb{R})$, its continuous dual is the space of **distributions** on $\mathbb{R}$. The space $C_c^\infty(\mathbb{R})$ is an LF-space, the direct limit of the spaces $C_c^\infty(\overline{\Omega})$, the subspace of $C^\infty(\mathbb{R})$ of functions with support in $\overline{\Omega}$. Each of these is a Fréchet space with norms

$$\|f\|_{\Omega,m} = \sup_\Omega \sum_{k<m} \left| d^k f/dx^k \right| .$$

By definition, if D is a distribution then for each relatively compact open subset Ω of $\mathbb{R}$, there exits m such that

$$\left|\langle D, f\rangle\right| \ll \|f\|_{\Omega,m} .$$

In other words, on any relatively compact open set a distribution has finite **order**.

For the moment, the most important distribution is the Dirac δ_0, where

$$\langle \delta_0, f\rangle = f(0) .$$

Any continuous function $F(x)$ on $\mathbb{R}$ determines a distribution by integration

$$\langle F, f\rangle = \int_{-\infty}^{\infty} F(x)f(x)\,dx\,.$$

If $F(x)$ is smooth and f lies in $C_c^\infty(\mathbb{R})$ then integration by parts tells us

$$\langle F', f\rangle = \int_{-\infty}^{\infty} F'(x)f(x)\,dx = -\int_{-\infty}^{\infty} F'(x)f(x)\,dx$$

so consistently with that formula the derivative of any distribution is defined by the formula

$$\langle D', f\rangle = -\langle D, f'\rangle\,.$$

For example,

$$\langle \delta_0', f\rangle = -\langle \delta_0, f'\rangle = -f'(0)$$

and $\langle \delta_0^{(m)}, f\rangle = (-1)^m f^{(m)}(0)$.

The distribution D is said to have **support** on the closed subset Ω of $\mathbb{R}$ if

$$\langle D, f\rangle = 0$$

whenever f lies in $C_c(\mathbb{R} - \Omega)$.

9.4. Proposition. *The Dirac distributions δ_0 and its derivatives span the subspace of distributions with support at the single point 0.*

Proof. Suppose Ω an interval around 0, m a bound on the order of D. It must be shown that D vanishes on the ideal of $C_c^\infty(\overline{\Omega})$ generated by x^m. This will follow if we can show that every function in (x^m) is approximated with respect to the norm $\|f\|_{\Omega,m}$ in $C_c^\infty(\overline{\Omega})$ by functions with support outside any given interval $|x| \ge \varepsilon$. For this choose a function χ which is 1 outside $|x| \ge 1$, vanishing inside the interval $|x| \le 1/2$, and monotonic in the intervals $[\pm 1, \pm 1/2]$. Then

$$\lim_{\varepsilon\to 0} \big\| f(x) - f(x)\chi(x/\varepsilon) \big\|_{[-1,1],m} = 0\,.$$

10. References

1. Nicholas Bourbaki, **Topological vector spaces**, Springer, translated from the French original dated 1981. This is one of the enthusiastic ones

2. Paul Garrett, 'Banach-Alaoglu, boundedness, weak-to-strong principles', preprint dated 2008:

`http://www/math.umn.edu/~garrett/`

3. Alexandre Grothendieck, **Topological vector spaces**, Gordon and Breach, 1973. For material that this book covers, it is often the clearest, with the most succinct proofs. It would be more useful if it had an index.

4. Michael Reed and Barry Simon, **Functional analysis**, volume I of the series *Methods of modern mathematical physics*, Academic Press, 1972. For the material that it covers, this is usually the best source. But it concentrates heavily on Hilbert spaces.

5. François Trèves, **Topological vector spaces, distributions, and kernels**, Academic Press, 1967. Invaluable, although sometimes the proofs of simple facts are very complicated.