Math 396. Connectedness of hyperplane complements

Note that the complement of a point in $\mathbf{R}$ is disconnected and the complement of a (translated) line in $\mathbf{R}^2$ is disconnected. Quite generally, we claim that the complement of a translated hyperplane in a finite-dimensional normed vector space over $\mathbf{R}$ is disconnected. In fact, this disconnectedness phenomenon is entirely an artifact of codimension 1. Higher codimensions never cause such problems. (Recall that if W is a subspace of a vector space V, the *codimension* of W in V is defined to be $\dim(V/W)$; this may be infinite, and it may be finite even if $\dim V$ and $\dim W$ are infinite, but when $\dim V$ is finite it is equal to $\dim V - \dim W$.)

As an example of the situation in higher codimension, we can see that removing a translated line (or several such) from $\mathbf{R}^3$ doesn't lead to disconnectedness: we can go "around" the line when thinking about paths linking up points. Likewise, removing the (codimension 2) origin from $\mathbf{R}^2$ doesn't cause disconnectedness. Roughly speaking, there's enough "elbow room" complementary to codimensions > 1 to avoid disconnectedness. The aim of this handout is to explore this situation.

1. Main result

We prove a connectivity theorem even when V is infinite-dimensional. Of course, to have a reasonable topology we suppose V is equipped with a norm, and we use the resulting metric topology (one can consider the possibility of putting a topology on V in other ways, but we will not discuss that here).

Theorem 1.1. *Let V be a normed vector space over $\mathbf{R}$. Then for any finite set of (not necessarily distinct) closed linear subspaces $W_1, \dots, W_k$ of (not necessarily finite) codimensions > 1 and any $v_1, \dots, v_k \in V$, the complement*

$$V - \bigcup_{i=1}^{k} (v_i + W_i)$$

is path-connected and hence is connected.

For example, the complement of any finite configuration of lines in $\mathbf{R}^3$ is path-connected; this is "geometrically obvious". Note that the theorem does not require V to be finite-dimensional, nor the W_i's to have finite codimension. However, we will construct the paths within well-chosen finite-dimensional subspaces of V, so the finite-dimensional case is the essential one. Of course, in the finite-dimensional case the closedness hypothesis on the subspaces is automatically satisfied. Also, infinitude of codimension of W_i is not a serious problem either: in fact, the higher the codimension of W_i gets, the more room there is complementary to W_i, and hence the easier life should be for finding paths! Before we prove the theorem, we record an interesting consequence.

Corollary 1.2. *Let V be a normed vector space over $\mathbf{C}$, and $W_1, \dots, W_k$ a finite collection of (not necessarily distinct) proper closed linear subspaces. For any $v_1, \dots, v_k \in V$, the complement*

$$V - \bigcup_{i=1}^{k} (v_i + W_i)$$

is path-connected and hence is connected.

Proof. We can view everything as $\mathbf{R}$-vector spaces at the expense of doubling dimensions and codimensions (when finite). In particular, V/W_j is a non-zero $\mathbf{C}$-vector space, whence as an $\mathbf{R}$-vector space has dimension at least 2 (perhaps infinite, which is even better!). Hence, by the theorem, we're done. ■

Now we give the proof of the theorem.

Proof. The case $\dim V \leq 1$ is trivial. Consider the special case $\dim V = 2$. In this case the only linear subspace of codimension > 1 is the origin, so we're just looking at the complement of a finite set of points. The path-connectedness of this is left to the reader as a pleasant exercise in using definitions.

Now consider the general case. Choose two points $x, y \in V$ not in any of the $v_i + W_i$'s. We want to find a path connecting them which avoids the complements. Translating everything in sight (i.e., x, y and the v_i's) by $-x$, we can assume $x = 0$ (so 0 is not in any $v_i + W_i$, so $v_i \notin W_i$ for all i). It is exactly the convenience of using such a translation (to reduce to studying paths joined to the origin) that forces us to formulate the original theorem in the context of translated subspaces. Since the linear subspaces W_i (and hence translates of them) are closed, the complement

$$V - \bigcup_{i=1}^{k} (v_i + W_i)$$

is *open.*

We first reduce to the finite-dimensional case. Since V/W_i has dimension at least 2, we get vectors $v'_{i,1}, v'_{i,2} \in V$ which induce linearly independent elements in V/W_i, which is to say

$$av'_{i,1} + bv'_{i,2} \notin W_i$$

for all $a, b \in \mathbf{R}$ not both zero. Let $\widetilde{V}$ be the *finite-dimensional* subspace of V spanned by $x = 0, y, v_1, \ldots, v_k$, and the vectors $v'_{i,1}, v'_{i,2}$ for $1 \leq i \leq k$, say given the induced norm from V. Let $\widetilde{W_i} = W_i \cap \widetilde{V}$. Since $av'_{i,1} + bv'_{i,2} \notin \widetilde{W_i}$ for all i, clearly $\widetilde{V}/\widetilde{W_i}$ has dimension at least 2 for all i (this is why we forced the $v'_{i,j}$'s to be in $\widetilde{V}$).

Because of all of the vectors we've forced into $\widetilde{V}$, it is easy to see that $\widetilde{V}$ and the $\widetilde{W_i}$'s with the v_i's satisfy all of the original hypotheses (especially the codimension > 1 condition). Hence, if we could settle the finite-dimensional case then we could find a continuous path in

$$\widetilde{V} - \bigcup (v_i + \widetilde{W_i}) = \widetilde{V} \bigcap \left(V - \bigcup_{i=1}^{k} (v_i + W_i) \right)$$

which joins $x = 0$ to y. Since $\widetilde{V} \to V$ is an isometry, this path is also continuous when viewed inside of V, and hence does the job.

Thus, we may now assume $\dim V < \infty$, and we will argue by induction on the dimension. Of course, in the finite-dimensional case all norms are equivalent and hence we can essentially suppress mention of the norm. As a preliminary step to help with the induction (basically to allow us to start the induction at dimension 2 rather than having to do an explicit argument in dimension 3 first), we reduce to the case where y is not in any of the W_i's.

We can find a small open ball $B_\varepsilon(y)$ around y which is inside of the complement of the closed $\cup(v_i + W_i)$, and even avoids touching any of the (finitely many, closed) W_i's which don't contain y. We claim there there is a $y' \in B_\varepsilon(y)$ not contained in any W_i's. Indeed, due to how we chose ε, we can use a translation by $-y$ to reduce to showing that inside of a *given* small ball around the origin we can always find a point which avoids any specified finite collection of hyperplanes. But any vector in V admits a non-zero scaling which is inside of $B_\varepsilon(0)$, so it is equivalent to show that V is not the union of finitely many hyperplanes. This follows from Lemma 2.3 below.

Using such a choice of $y' \in B_\varepsilon(y)$, if we can find a path from 0 to y' in the complement of the $(v_i + W_i)$'s, then hooking this onto a radial path from y' to y in the ball $B_\varepsilon(y)$ (which is likewise disjoint from all $(v_i + W_i)$'s), we'll be done. Hence, replacing y with a well-chosen sufficiently close y' lets us assume that $y \notin W_i$ for all i.

Now the idea is to take a suitably well-chosen hyperplane slice through y to drop the dimension of V without affecting codimensions of the W_i's. This will reduce us to the case $\dim V = 2$ which has already been treated. More specifically, we will find a 2-dimensional subspace V_0 in V which contains y but meets each W_i in $\{0\}$ (this would not be possible if $y \in W_i$!).

Now quite generally, if $U, U' \subseteq V$ are linear subspaces and $v, v' \in V$ are points, then it is easy to see that

$$(v + U) \cap (v' + U') = \begin{cases} \emptyset, & \text{if } v - v' \notin U + U' \\ (v' + u') + (U \cap U'), & \text{if } v - v' = u + u' \in U + U' \end{cases}$$

Thus, back in our original situation, if V_0 is a 2-dimensional subspace of V which contains y and meets each W_i in $\{0\}$ then $V_0 \cap (v_i + W_i)$ is either empty or a point. Thus, we have

$$V_0 \cap (V - \bigcup_{i=1}^{k} (v_i + W_i)) = V_0 - \bigcup_{i=1}^{k} (V_0 \cap (v_i + W_i))$$

with each $V_0 \cap (v_i + W_i)$ either empty or a point. Thus, slicing with the subspace V_0 which contains 0 and y brings us to a complement of a finite set in the 2-dimensional V_0, and such a complement is path-connected (and contains y and $0 = x$). Our problem is now reduced to a statement in linear algebra which we can prove over an arbitrary *infinite* field, as in the theorem below (in which the "auxiliary vector" is y). The required result in linear algebra is treated in the next section. ■

2. A theorem in linear algebra

To emphasize the essentially algebraic (as opposed to topological) aspect of what remains to be done, we now work over an essentially arbitrary field.

Theorem 2.1. *Let V be a finite-dimensional vector space over an infinite field F, with $\dim V \geq 2$, and let $W_1, \ldots, W_k$ be (not necessarily distinct) linear subspaces of codimensions > 1. Choose an auxiliary vector $v_0 \in V$ with $v_0 \notin W_i$ for all i. Then there exists a 2-dimensional subspace V_0 such that $v_0 \in V_0$ and $V_0 \cap W_i = \{0\}$ for all i.*

Remark 2.2. This lemma is *false* over finite fields. Indeed, over a finite field a finite-dimensional vector space V contains only finitely many vectors, so we can even find finitely many *lines* (e.g., the span of each non-zero element) whose union is the entire space. Taking $\{W_i\}$ to be the finite set of lines in V, any non-zero subspace certainly contains one of these lines and so no such V_0 as in the theorem can exist. It is a characteristic of infinite fields that a vector space of finite dimension > 1 over such a field cannot be expressed as a finite union of lower-dimensional subspaces, and we will reduce the proof of the theorem to exactly this fact, which will be proven afterwards as a separate lemma (that was already used in earlier arguments).

Proof. We can drop any W_i's which are equal to 0, so we may assume $W_i \neq 0$ for all i (and that there actually are some W_i's). We induct on $\dim V \geq 2$, the case $\dim V = 2$ being clear (as then the W_i's are automatically $\{0\}$, so we may use $V_0 = V$). When $\dim V > 2$, we just need to find a codimension-1 subspace H such that $v_0 \in H$ and $W_i \not\subseteq H$ for all i. Indeed, in that case $W_i + H = V$ (as W_i then surjects onto the 1-dimensional V/H), so

$$\begin{aligned} \dim(W_i \cap H) &= \dim(W_i) + \dim(H) - \dim(W_i + H) \\ &= \dim(W_i) + \dim(H) - \dim(V) \\ &= \dim(H) - \dim(V/W_i) \end{aligned}$$

In other words, if we slice with a hyperplane H not containing W_i, then the codimension

$$\dim H/(W_i \cap H) = \dim(H) - \dim(W_i \cap H) = \dim V/W_i$$

of $W_i \cap H$ in H is equal to the codimension of W_i in V, which we assumed to be > 1.

Thus, once we find an H containing v_0 which does not contain any of the W_i's, then we can replace V with H and each W_i with $W_i \cap H$ without destroying any of the hypotheses (and any $W_i \cap H$'s which vanish can be dropped). Since $\dim H = \dim V - 1$, by induction on the dimension of V we'd be done. Our problem therefore is to find a codimension-1 subspace through v_0 which does not contain any of the *non-zero* codimension subspaces $W_1, \dots, W_k$ whose codimension in V is > 1.

As a concrete example, for $V = \mathbf{R}^3$ this says that, given any finite set of lines $L_1, \dots, L_k$ in $\mathbf{R}^3$ and any point v_0 *not* on any of these lines, we can find a plane through v_0 which does not contain any of the lines. It is geometrically obvious in this case that a "random" choice of plane through v_0 will do the job (though a few "bad" planes may fail).

The general argument goes as follows. We can view the problem of constructing a hyperplane H as the problem of finding a suitable non-zero linear functional $\ell : V \to F$ (with H then taken to be the codimension-1 kernel of ℓ). In such terms, we seek a point ℓ in the dual space $V^\vee$ with $\ell(v_0) = 0$ but ℓ non-zero on each of the *non-zero* subspaces W_i. This ensures that $H = \ker \ell$ is a hyperplane containing v_0 but not any of the W_i's. Consider the annihilator $W_i^\perp \subseteq V^\vee$, which is to say the subspace of functionals which vanish on W_i. Since linear maps $V \to V'$ that kill a subspace W uniquely factor through the projection $V \to V/W$, by taking the case $V' = F$ we arrive at an evident linear isomorphism

$$W_i^\perp \simeq (V/W_i)^\vee,$$

so this subspace of $V^\vee$ has dimension $\dim V/W_i < \dim V = \dim V^\vee$, and hence it is a proper subspace of $V^\vee$ (with codimension $\dim W_i$).

Let $K = (Fv_0)^\perp$, a codimension-1 subspace of $V^\vee$. Since $v_0 \notin W_i$, we have $Fv_0 \not\subseteq W_i$, so $W_i^\perp \not\subseteq K$ (as otherwise applying $(\cdot)^\perp$ to this via $V \simeq V^{\vee\vee}$ would yield the reverse inclusion $Fv_0 \subseteq W_i$ which we have assumed not to hold). For each i, we claim that the subspace $K_i = K \cap W_i^\perp$ in K is a *proper* subspace. If not, so this intersection equals K, then we'd get $K \subseteq W_i^\perp \subsetneq V^\vee$, forcing $W_i^\perp = K$ since K has codimension 1 in $V^\vee$ (so there are no non-trivial intermediate subspaces between K and $V^\vee$). But we've just seen that $W_i^\perp$ is not contained in K, so this would be a contradiction.

Now the task is to show that the vector space $K = (Fv_0)^\perp$ inside of $V^\vee$ contains an element which is *not* inside of any of the proper subspaces $K_i = W_i^\perp \cap K$. In other words, we want to show that the vector space K cannot be a union of finitely many proper subspaces (which would be false over a finite field). So far we have not used that F is an infinite field, but it is at this step that the infinitude of F plays the crucial role. We isolate the necessary fact in the form of a lemma below. ∎

Lemma 2.3. *Let F be an infinite field, V a vector space, and $V_1, \dots, V_k$ finitely many proper subspaces. Then V is not the union of the V_j's.*

Proof. The cases $k = 0, 1$ are clear. This also settles $\dim V \le 1$. The idea now is to draw a "random" line in V and to find a point on this line which is not on any of the V_i's.

We may assume $k > 1$ and (by induction on k) the result is known for collections of $< k$ proper subspaces. By induction we can choose a vector v not contained in $V_1, \dots, V_{k-1}$. If also $v \notin V_k$, we're done. Otherwise, choose another vector v' not contained in the proper subspace V_k (so $v' \neq v$). Let L be the span of $v' - v \neq 0$, so the translated line $v + L = v' + L$ passes through both v' and

v. Note that $L \cap V_i = \{0\}$ since this intersection is a *proper* subspace of the 1-dimensional L (as L contains both v and v', at least one of which is not in V_i).

Consider the intersection

$$(v + L) \cap V_i = (v' + L) \cap V_i$$

for $1 \leq i \leq k$. Since $L \cap V_i = \{0\}$, this intersection $(v + L) \cap V_i$ is either empty or a point (i.e., it cannot contain two points, as the difference would be a non-zero element in $L \cap V_i = \{0\}$). Thus,

$$(v + L) \cap (\bigcup_{i=1}^{k} V_i) = \bigcup_{i=1}^{k} ((v + L) \cap V_i)$$

is a *finite* (perhaps empty) union of points. But $v + L$ is *infinite* since L is 1-dimensional over an infinite field. Hence, we can find $x \in v + L$ not contained in any V_i. The union of the V_i's is therefore not all of V. ∎