INFINITE-DIMENSIONAL DUAL SPACES

KEITH CONRAD

Let K be a field and V be a K-vector space. The dual space $V^\vee$ is the set of all K-linear functions $\varphi\colon V \to K$, with a natural concept of addition and K-scaling of functions with values in K. If $\dim_K(V)$ is finite then $\dim_K(V) = \dim_K(V^\vee)$, so V and $V^\vee$ are isomorphic K-vector spaces (non-canonically) and there is a canonical isomorphism from V to its double dual $V^{\vee\vee}$ by $v \mapsto [\varphi \mapsto \varphi(v)]$. If $\dim_K(V)$ is infinite, however, there is no isomorphism between V and $V^{\vee\vee}$ because their dimensions are not equal by the following result.

Theorem. *If V is an infinite-dimensional K-vector space then* $\dim_K(V) < \dim_K(V^\vee)$.

By replacing V with $V^\vee$ in this theorem, we get $\dim_K(V) < \dim_K(V^\vee) < \dim_K(V^{\vee\vee})$ when V is infinite-dimensional over K, so $V \not\cong V^{\vee\vee}$ as K-vector spaces.

It is important to realize that this theorem uses the algebraic notion of dimension as the cardinality of a linearly independent spanning set, where the allowed linear combinations use *finitely many* nonzero vectors at a time. This is a contrast to vector spaces over **R** or **C** having a notion of convergence (topological vector spaces), such as Banach or Hilbert spaces, where linear combinations with infinitely many nonzero terms occur. The meaning of dimension using such linear combinations is *more restrictive* than the algebraic meaning of dimension since the meaning of a basis changes (linear dependence relations can have infinitely many nonzero terms) and the meaning of the dual space also changes: it is the space of all *continuous* linear functionals to **R** or **C**, not all linear functionals to **R** or **C**.

For example, each infinite-dimensional Hilbert space (*e.g.*, ℓ^2 or $L^2([0,1])$) is isomorphic to its (continuous) dual space and to its double (continuous) dual space. That does not contradict the theorem above, since the meaning of dimension and dual space in analysis is not the algebraic meaning of those terms: the analyst's dimension and dual space are smaller than the algebraist's dimension and dual space when the dimension is infinite.

Proof. This argument is based on an answer by Andrea Ferretti on MathOverflow [1].

Let $\{e_i\}_{i\in I}$ be a K-basis of V, so I is infinite and $V = \bigoplus_{i\in I} Ke_i$. Each member of $V^\vee$ is determined by its values on the basis. For every collection of scalars $c_i \in K$ ($i \in I$) there is a unique K-linear functional $\varphi \in V^\vee$ where $\varphi(e_i) = c_i$ for all $i \in I$, and this gives an isomorphism $V^\vee \cong \prod_{i\in I} K$ by $\varphi \mapsto (\varphi(e_i))_{i\in I}$. Therefore we can view $V^\vee$ as all functions $I \to K$ and V is the functions $I \to K$ with *finite support* (vanishing outside a finite subset). By definition, $\dim_K(V) = \operatorname{card}(I)$. It is plausible that $\prod_{i\in I} K$ has a larger dimension than $\bigoplus_{i\in I} K$ when I is infinite, and this is what we will prove.

Step 1: $\dim_K(V) \le \dim_K(V^\vee)$.

To each e_i we can associate the functional $\delta_{e_i} \in V^\vee$ that is 1 at e_i and 0 at elements of the basis besides e_i. Linear independence of the e_i in V implies linear independence of the δ_{e_i} in $V^\vee$, so $\dim_K(V) = \operatorname{card}(I) \le \dim_K(V^\vee)$.

Step 2: If the field K is finite or countably infinite then $\operatorname{card}(V) = \operatorname{card}(I)$.

Note the left side is $\operatorname{card}(V)$: the cardinality of the whole set V and not just the cardinality of a basis of V, which is $\operatorname{card}(I)$.

Looking at coefficients of elements of V in its basis $\{e_i : i \in I\}$ lets us think of V as the set of functions $I \to K$ with finite support. For $n \geq 1$, let $S_n(I)$ be the n-element subsets of I. By ordering each n-element subset of I in some way, $S_n(I)$ can be viewed (non-canonically) as a subset of I^n. Since I is infinite, $\text{card}(I^n) = \text{card}(I)$. Thus $\text{card}(S_n(I)) \leq \text{card}(I^n) = \text{card}(I)$. (In fact, $\text{card}(S_n(I)) = \text{card}(I)$, but we won't use this.)

The functions $I \to K$ with support of size n can be viewed (non-canonically) as a subset of $S_n(I) \times K^n$. Since $\text{card}(S_n(I)) \leq \text{card}(I)$ and K is at most countable, $\text{card}(S_n(I) \times K^n) \leq \text{card}(I)$. Thus the set of functions $I \to K$ with support of size n has cardinality at most $\text{card}(I)$, and taking a countable union tells us the set of functions $I \to K$ with finite support has cardinality at most $\text{card}(I)$, so $\text{card}(V) \leq \text{card}(I)$.

Inside V, the basis $\{e_i : i \in I\}$ tells us $\text{card}(I) \leq \text{card}(V)$. From $\text{card}(V) \leq \text{card}(I)$ and $\text{card}(I) \leq \text{card}(V)$, we get $\text{card}(V) = \text{card}(I)$ by the Schroeder–Bernstein theorem.

Step 3: If the field K is finite or countably infinite then $\dim_K(V) < \dim_K(V^\vee)$.

Since $V^\vee$ is the set of *all* functions $I \to K$ and $\{0, 1\} \subset K$, $\text{card}(V^\vee)$ is at least as large as the cardinality of the set of all functions $I \to \{0, 1\}$. Such functions naturally correspond to the subsets of I, so $\text{card}(V^\vee) \geq \text{card}(2^I) > \text{card}(I)$, where the strict inequality follows from Cantor's diagonalization argument. Thus $\text{card}(V^\vee) > \text{card}(I)$. We have $\text{card}(I) = \text{card}(V)$ by Step 2, so $\text{card}(V^\vee) > \text{card}(V)$. Therefore $V \not\cong V^\vee$ since isomorphic vector spaces have the same cardinality. A K-vector space is determined up to isomorphism by its dimension over K, so $\dim_K(V) \neq \dim_K(V^\vee)$. Since $\dim_K(V) \leq \dim_K(V^\vee)$ by Step 1, we have $\dim_K(V) < \dim_K(V^\vee)$.

Step 4: If the field K is arbitrary then $\dim_K(V) < \dim_K(V^\vee)$.

Inside the field K is a finite or countably infinite subfield, namely its (unique) minimal subfield

$$F \cong \begin{cases} \mathbf{Q}, & \text{if } K \text{ has characteristic } 0, \\ \mathbf{F}_p, & \text{if } K \text{ has characteristic } p. \end{cases}$$

Let W be the F-span of the K-basis e_i of V:

$$V = \bigoplus_{i \in I} Ke_i, \quad W = \bigoplus_{i \in I} Fe_i \subset V.$$

By definition, $\dim_F(W) = \text{card}(I) = \dim_K(V)$. Using Step 3 with F in place of K tells us $\dim_F(W) < \dim_F(W^\vee)$, where $W^\vee$ is the F-dual space of W. We will show $\dim_F(W^\vee) \leq \dim_K(V^\vee)$, so

$$\dim_K(V) = \dim_F(W) < \dim_F(W^\vee) \leq \dim_K(V^\vee),$$

and thus $\dim_K(V) < \dim_K(V^\vee)$. Our proof is reduced to showing $\dim_F(W^\vee) \leq \dim_K(V^\vee)$.

We construct an F-linear mapping $W^\vee \to V^\vee$. For $\varphi \in W^\vee$, meaning $\varphi \colon W \to F$ is F-linear, define $\widetilde{\varphi} : V \to K$ as follows: for $v \in V$ we can write v uniquely as $\sum a_i e_i$ with $a_i \in K$ where all but finitely many a_i are 0. Set $\widetilde{\varphi}(v) = \sum a_i \varphi(e_i)$, which *makes sense* since it is a finite sum in K. This construction $\varphi \rightsquigarrow \widetilde{\varphi}$ has the following properties:

(1) Each function $\widetilde{\varphi} \colon V \to K$ is K-linear: writing $v' = \sum a_i' e_i$ with $a_i' \in K$ and picking α in K,

$$\widetilde{\varphi}(v + v') = \sum (a_i + a_i')\varphi(e_i) = \widetilde{\varphi}(v) + \widetilde{\varphi}(v'), \quad \widetilde{\varphi}(\alpha v) = \sum (\alpha a_i)\varphi(e_i) = \alpha\widetilde{\varphi}(v).$$

(2) For $w \in W$ we have $\widetilde{\varphi}(w) = \varphi(w)$: write $w = \sum a_i e_i$ with $a_i \in F$, so by F-linearity of φ,

$$\widetilde{\varphi}(w) = \sum a_i \varphi(e_i) = \varphi\left(\sum a_i e_i\right) = \varphi(w).$$

(3) The mapping $W^\vee \to V^\vee$ given by $\varphi \mapsto \widetilde{\varphi}$, from an F-dual space to a K-dual space, is F-linear: $\widetilde{\varphi + \psi} = \widetilde{\varphi} + \widetilde{\psi}$ and $\widetilde{c\varphi} = c\widetilde{\varphi}$ for $c \in F$. Details are left to the reader.

Property (2) tells us that $\widetilde{\varphi}$ remembers φ, since $\widetilde{\varphi}|_W = \varphi$, so the mapping $W^\vee \to V^\vee$ given by $\varphi \mapsto \widetilde{\varphi}$ is injective. Therefore we can view $W^\vee$ as an F-linear subspace of $V^\vee$, which implies $\dim_F(W^\vee) \le \dim_F(V^\vee)$. Our goal was to show $\dim_F(W^\vee) \le \dim_K(V^\vee)$, which is stronger[1]. Here is a situation where that stronger inequality is valid.

Claim: Let F be a subfield of K and U be a K-vector space with an F-linear subspace U'. If every finite F-linearly independent subset of U' is K-linearly independent, then $\dim_F(U') \le \dim_K(U)$.

Proof of Claim: We'll prove the contrapositive. Assume $\dim_F(U') > \dim_K(U)$. Let $\{u'_j\}_{j\in J}$ be an F-basis of U'. It is a subset of U, so the dimension inequality tells us that $\{u'_j\}_{j\in J}$ is not K-linearly independent: there is some *finite* nontrivial (= not all coefficients are 0) K-linear combination from this F-basis that is 0. That gives us a finite F-linearly independent subset of U' that is not K-linearly independent, which proves (the contrapositive of) the claim.

We will apply this claim to $U = V^\vee$ and $U' = W^\vee$, where $W^\vee$ is embedded into $V^\vee$ as the mappings $\widetilde{\varphi}$ for $\varphi \in W^\vee$. To use the claim, we will show for each finite F-linearly independent subset $\{\varphi_1, \ldots, \varphi_n\}$ of $W^\vee$ that the set $\{\widetilde{\varphi}_1, \ldots, \widetilde{\varphi}_n\}$ is K-linearly independent in $V^\vee$.

Suppose $\sum_{j=1}^n c_j\widetilde{\varphi}_j = 0$ in $V^\vee$, where $c_j \in K$. We want to show $c_1, \ldots, c_n$ are all 0. For every $w \in W$,

$$0 = \sum_{j=1}^n c_j\widetilde{\varphi}_j(w) = \sum_{j=1}^n c_j\varphi_j(w). \tag{1}$$

If, for each k from 1 to n, there is a $w_k \in W$ such that $\varphi_k(w_k) = 1$ and $\varphi_j(w_k) = 0$ if $j \ne k$, then using this w_k for w in (1) tells us $c_k = 0$, so all the coefficients $c_1, \ldots, c_n$ are 0 and we'd be done. Finding the w_k is entirely a problem about linear algebra in an F-vector space W and its F-dual space $W^\vee$.

To prove such w_k exist, we'll show the mapping $W \to F^n$ where $w \mapsto (\varphi_1(w), \ldots, \varphi_n(w))$ is surjective. This mapping is F-linear since $\varphi_1, \ldots, \varphi_n$ are each F-linear, so the image is a subspace of F^n. The ordinary dot product on F^n is a non-degenerate bilinear form (we use non-degeneracy because positive-definiteness would make no sense if $F = \mathbf{F}_p$), so to show the image of the mapping $W \to F^n$ is all of F^n it suffices to show the only element of F^n with dot product 0 against all elements in the image is the zero vector. Let $(a_1, \ldots, a_n) \in F^n$ have dot product 0 with the image, so for all $w \in W$ we have

$$(a_1, \ldots, a_n) \cdot (\varphi_1(w), \ldots, \varphi_n(w)) = 0$$

or equivalently $\sum_{j=1}^n a_j\varphi_j(w) = 0$. This being true for all $w \in W$ makes $\sum_{j=1}^n a_j\varphi_j = 0$ in $W^\vee$, so $a_1, \ldots, a_n$ are all 0 because $\{\varphi_1, \ldots, \varphi_n\}$ is an F-linearly independent subset of W. $\square$

References

[1] Andrea Ferreti, answer to MathOverflow question 13322, Slick proof?: A vector space has the same dimension as its dual if and only if it is finite dimensional, http://mathoverflow.net/questions/13322.

[1] This is stronger since $\dim_K(V^\vee) \le \dim_F(V^\vee)$: K-linear independence implies F-linear independence.