

BASES FOR INFINITE DIMENSIONAL VECTOR SPACES
MATH 513 LINEAR ALGEBRA SUPPLEMENT

PROFESSOR KAREN E. SMITH

We have proven that every finitely generated vector space has a basis. But what about vector spaces that are not finitely generated, such as the space of all continuous real valued functions on the interval $[0, 1]$? Does such a vector space have a basis?

By definition, a basis for a vector space V is a linearly independent set which generates V . But we must be careful what we mean by linear combinations from an infinite set of vectors. The definition of a vector space gives us a rule for adding two vectors, but not for adding together infinitely many vectors. By successive additions, such as $(v_1 + v_2) + v_3$, it makes sense to add any *finite* set of vectors, but in general, there is no way to ascribe meaning to an infinite sum of vectors in a vector space.

Therefore, when we say that a vector space V is *generated by* or *spanned by* an infinite set of vectors $\{v_1, v_2, \dots\}$, we mean that each vector v in V is a *finite* linear combination $\lambda_{i_1}v_{i_1} + \dots + \lambda_{i_n}v_{i_n}$ of the v_i 's. Likewise, an infinite set of vectors $\{v_1, v_2, \dots\}$ is said to be *linearly independent* if the only *finite* linear combination of the v_i 's that is zero is the trivial linear combination. So a set $\{v_1, v_2, v_3, \dots\}$ is a basis for V if and only if every element of V can be written in a *unique* way as a *finite* linear combination of elements from the set.

Actually, the notation $\{v_1, v_2, v_3, \dots\}$ for an infinite set is misleading because it seems to indicate that the set is *countable*. We want to allow the possibility that a vector space may have an *uncountable* basis. For this reason, it is more sensible to use notation such as $\{v_\lambda | \lambda \in \Lambda\}$, where Λ is some unspecified indexing set.

Does every vector space have a basis? Let's see what happens in the simplest cases.

Example 1. Consider the vector space $\mathbb{P}(\mathbb{R})$ of all polynomial functions on the real line. This vector space is not generated by any finite set. On the other hand, every polynomial is a finite linear combination of the polynomials

$$f_n(x) = x^n$$

for $n = 0, 1, 2, \dots$, so these polynomials span $\mathbb{P}(\mathbb{R})$. Furthermore, if a polynomial $\lambda_n x^n + \lambda_{n-1} x^{n-1} + \dots + \lambda_1 x + \lambda_0$ is the zero function on \mathbb{R} , then all coefficients λ_i

must be zero, so the f_n 's are linearly independent. In other words, the functions f_n form a basis for the vector space $\mathbb{P}(\mathbb{R})$.

Example 2. Let \mathbb{R}^∞ be the vector space of infinite sequences $(\alpha_1, \alpha_2, \alpha_3, \dots)$ of real numbers. This is a natural generalization of \mathbb{R}^n . The vector addition and scalar multiplication are defined in the natural way: the sum of $(\alpha_1, \alpha_2, \alpha_3, \dots)$ and $(\beta_1, \beta_2, \beta_3, \dots)$ is $(\alpha_1 + \beta_1, \alpha_2 + \beta_2, \alpha_3 + \beta_3, \dots)$ and the product of $(\alpha_1, \alpha_2, \alpha_3, \dots)$ by a scalar λ is the sequence $(\lambda\alpha_1, \lambda\alpha_2, \lambda\alpha_3, \dots)$. What might be a basis for \mathbb{R}^∞ ?

A natural set to consider is the set $\{e_1, e_2, \dots\}$ generalizing the standard basis for \mathbb{R}^n . In other words, e_i here is the sequence of all zeroes except in the i -th spot where there appears a 1. The vectors $\{e_i\}$ are obviously linearly independent. But do they span \mathbb{R}^∞ ? Note that every finite linear combination of the e_i 's will be a vector in which only finitely many components are non-zero. Thus there is no way to write vector like $(1, 1, 1, \dots)$ as a finite linear combination of the e_i 's.

Since the vector $v = (1, 1, 1, \dots)$ can not be written as a linear combination of the vectors $e_i = (0, \dots, 0, 1, 0, 0, \dots)$, it must be that together the e_i 's and v form a linearly independent set. If this set spans \mathbb{R}^∞ , it must be a basis. Unfortunately, it does not span: for example, the vector $w = (1, 2, 3, 4, \dots)$ is not a linear combination of the e_i 's and v . We could try to repeat this process, checking whether the set $\{w, v, e_1, e_2, e_3, \dots\}$ spans \mathbb{R}^∞ , and if not, enlarging the set to a bigger linearly independent set. But does the process eventually end, yielding a basis for \mathbb{R}^∞ ? Actually, it can be proved that no countable set of vectors in \mathbb{R}^∞ spans \mathbb{R}^∞ , so even if \mathbb{R}^∞ has a basis, we will not be able to construct it by simply adding in elements one at a time to this set.

One might try to argue that the vectors e_i generate \mathbb{R}^∞ if we allow "infinite linear combinations" of them. Indeed, we could think of $(\alpha_1, \alpha_2, \alpha_3, \dots)$ as the sum $\sum_{i=1}^{\infty} \alpha_i e_i$. This is unsatisfactory for several reasons. First, although this particular infinite sum seems to make sense, infinite sums in general do not. For example, what would be the sum $(1, 1, 1, \dots) + (2, 2, 2, \dots) + (3, 3, 3, \dots)$? This could become a serious problem if we try to do any kind of algebraic manipulations with these "infinite sums". Also, the fact that we just happen to make sense of an infinite sum of a special type of vectors in a special type of vector space does not generalize well to other settings. Even for subspaces of \mathbb{R}^∞ , these "infinite sums" may not make sense.

Example 3. Let V be the subspace of \mathbb{R}^∞ consisting of those sequences $(\alpha_1, \alpha_2, \dots)$ for which the limit $\lim a_n$ exists as n goes to infinity. That is, V is the space of convergent sequences of real numbers. Clearly the elements e_i above are in V , since the limit of the sequence $(0, \dots, 0, 1, 0, 0, \dots)$ is zero. Of course, the e_i are still linearly independent. Do they span V ? No, again the element $(1, 1, 1, \dots)$ is in V (the limit of this sequence is 1), but not in the span of the e_i . It is not clear whether we can add in some more vectors of V so as to enlarge this set to a basis.

Note that we cannot describe V as the set of "infinite linear combinations" of the e_i . Indeed, the set of all the "infinite linear combinations" of the e_i is all of \mathbb{R}^∞ , whereas not every vector in \mathbb{R}^∞ is a convergent sequence. For example

$(1, 2, 3, 4, \dots)$ is not in V .

Example 4. Let ℓ^2 be the subspace of \mathbb{R}^∞ of all square summable sequences, that is, of vectors $(\alpha_1, \alpha_2, \dots)$ such that $\sum_{i=1}^\infty \alpha_i^2$ exists. This space includes the linearly independent vectors e_i , but is not spanned by the e_i . For example, the vector $(1, \frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}, \dots)$ is in ℓ^2 (it has square-sum $4/3$), but is not a (finite) linear combination of the e_i . Does this space have a basis?

Example 5. Let V be the subspace of \mathbb{R}^∞ spanned by the e_i for $i = 1, 2, 3, \dots$. Since the e_i are linearly independent, clearly they form a basis for V . What is the space V ? It consists of all vectors in \mathbb{R}^∞ in which all but finitely many of the slots are zeroes. This can be interpreted as the union of all \mathbb{R}^n as n goes to infinity.

Example 6. Let \mathbb{R} be the set of real numbers considered as a vector space over the field \mathbb{Q} of rational numbers. What could possibly be a basis? The elements $\sqrt{2}, \sqrt{3}, \sqrt{5}, \sqrt{6}, \dots$, can be shown to be linearly independent, but they certainly don't span \mathbb{R} , as we also need elements like π, π^2, π^3, \dots , which also form a linearly independent set. In fact, because \mathbb{Q} is countable, one can show that the subspace of \mathbb{R} generated by any countable subset of \mathbb{R} must be countable. Because \mathbb{R} itself is uncountable, no countable set can be a basis for \mathbb{R} over \mathbb{Q} . This means that any basis for \mathbb{R} over \mathbb{Q} , if one exists, is going to be difficult to describe.

These examples make it clear that even if we *could* show that every vector space has a basis, it is unlikely that a basis will be easy to find or to describe in general.

Every vector space has a basis.

Although it may seem doubtful after looking at the examples above, it is indeed true that *every vector space has a basis*. Let us try to prove this. First, consider any linearly independent subset of a vector space V , for example, a set consisting of a single non-zero vector will do. Call this set \mathcal{S}_1 . If \mathcal{S}_1 spans V , it is a basis, and the proof is complete. If not, we can *choose* a vector of V not in \mathcal{S} and the union $\mathcal{S}_2 = \mathcal{S}_1 \cup \{v\}$ is a larger linearly independent set. Either this set is a basis, or we can again enlarge it by choosing some vector of V not in the span. We can repeat this process over and over, and hope that it eventually ends. But it is easy to see that such a naive approach will not work in general unless V is finite dimensional. Indeed, starting from \mathcal{S}_1 being a single element set, every \mathcal{S}_i produced this way will be finite.

On the other hand, using this idea, we either get a basis for V eventually or we get an increasing collection of linearly independent sets

$$\mathcal{S}_1 \subset \mathcal{S}_2 \subset \mathcal{S}_3 \subset \dots$$

Clearly the union \mathcal{S} of all the \mathcal{S}_i is a linearly independent set, since any *finite* linear combination of the elements of the union must involve elements from one of the sets \mathcal{S}_i . If this set \mathcal{S} spans V , it is a basis and we are done. This is what happens, for example, if V is the space of polynomials defined in Example 1 and \mathcal{S}_i is the set

$\{1, x, x^2, \dots, x^i\}$. But in general, the set \mathcal{S} may not span V . For example, if V is the vector space \mathbb{R}^∞ , the set \mathcal{S}_i could be $\{e_1, e_2, \dots, e_i\}$. The union of the sets \mathcal{S}_i produces the set \mathcal{S} of all e_i , which as we have seen above, does not span \mathbb{R}^∞ .

However, even if \mathcal{S} does not span V , it is at least linearly independent, so we could again choose a vector v not in the span of \mathcal{S} . By adding v to \mathcal{S} , we again get a larger linearly independent set, and we can repeat the process. Does this process eventually terminate, producing for us a basis of V ? This is not at all clear.

To prove that every vector space has a basis, we need *Zorn's Lemma*. Assume \mathfrak{C} is a collection of subsets of some fixed unnamed set, and assume that \mathfrak{C} has the property that whenever there is a chain $\mathcal{S}_1 \subset \mathcal{S}_2 \subset \dots$ of sets in \mathfrak{C} , the union of this chain also belongs to \mathfrak{C} . Then Zorn's Lemma says that \mathfrak{C} contains a maximal element. This means that \mathfrak{C} contains some set \mathcal{M} which is not properly contained in any other set in the collection \mathfrak{C} . In fact, Zorn's lemma implies that *every* set \mathcal{S} in \mathfrak{C} is contained in some maximal set \mathcal{M} , because we can apply Zorn's lemma to the subcollection of sets in \mathfrak{C} containing \mathcal{S} .

Now let \mathfrak{C} be the collection of all linearly independent subsets of a vector space V . Since the union of any increasing chain

$$\mathcal{S}_1 \subset \mathcal{S}_2 \subset \dots$$

of linearly independent sets is also a linearly independent set, Zorn's Lemma implies that there is a maximal linearly independent set \mathcal{M} . This maximal linearly independent set is a basis for V . Indeed, if it doesn't span V , we could choose a vector in V but not in the span of \mathcal{M} , and by adding it to \mathcal{M} , we could enlarge our supposedly maximal linearly independent set. This contradiction completes the proof that every vector space V has a basis. Zorn's Lemma also implies that every linearly independent set can be extended to a basis, because given any linearly independent set \mathcal{S} , we know that there is some maximal linearly independent set containing \mathcal{S} .

There is a major drawback to this proof that every vector space has a basis: unless the dimension is finite, or at least countable, it doesn't give us any idea how to actually find a basis. In fact, this is a serious problem with the concept of a basis for infinite dimensional spaces in general. Although Zorn's Lemma tells us a basis exists, in practice, this fact may be useless if we do not have a procedure for finding one. For this reason, bases are not often used for infinite dimensional spaces, and mathematicians have come up with alternative ideas. A basis for an infinite dimensional space is also called a *Hamel basis* to distinguish it from different but more useful notions that will also be (confusingly!) called bases.

So how does one prove Zorn's Lemma? Although we won't prove it here, Zorn's Lemma is logically equivalent to the "axiom of choice". The axiom of choice says that given any collection \mathfrak{C} of sets, we can choose an element x from each set \mathcal{S} of \mathfrak{C} . This may seem "obvious"—or does it? There is of course no problem if there

are finitely many sets in the collection, but what if there are infinitely many, maybe even uncountably many? The axiom of choice and Zorn's Lemma bothered many mathematicians (and still bothers some!) for various reasons. For example, using the axiom of choice, one can prove that a ball the size of the sun can be cut into finitely many pieces and then reassembled into a ball the size of a pinhead. So if we accept the axiom of choice (and equivalently, Zorn's Lemma), we must accept such statements.

There is no way to prove the axiom of choice: one either accepts it as an axiom or one doesn't. The axiom of choice (and so the equivalent formulation Zorn's Lemma) is logically independent from the other axioms of set theory, a fact proven by Paul Cohen in 1963. In other words, we derive no contradictions if we assume it is true, and we derive no contradictions if we assume it is false. The axiom of choice is no longer as controversial as it once was. It is accepted by most mathematicians these days, but the degree to which it is used without comment depends on the branch of mathematics. To learn more about the axiom of choice and related matters, you can take a course on mathematical logic. Many books on analysis also contain discussions of the axiom of choice, for example, [HS] has nice discussion, including the proof of the equivalence of Zorn's Lemma and the axiom of choice.

Summary: Every vector space has a basis, that is, a maximal linearly independent subset. Every vector in a vector space can be written in a unique way as a finite linear combination of the elements in this basis. A basis for an infinite dimensional vector space is also called a *Hamel basis*. The proof that every vector space has a basis uses the axiom of choice. There is no practical way to find a Hamel basis in general, which means we have little use for the concept of a basis for a general infinite (especially uncountable) dimensional vector space.

Alternatives to the concept of a basis.

There are some vector spaces, such as \mathbb{R}^∞ , where at least certain infinite sums make sense, and where every vector can be uniquely represented as an *infinite* linear combination of vectors. Although there are some drawbacks to this idea, it is also very useful to use the idea of infinite sums in some cases, provided we can make sense out of them.

What do we need in order to make sense of an infinite sum of vectors? The only way to make sense out of an infinite sum $\sum_{i=1}^{\infty} v_i$ is to consider the sequence of partial sums $s_n = \sum_{i=1}^n v_i$ (for $n = 1, 2, 3, \dots$) and hope that this sequence $\{s_n\}$ converges. Thus we must have a notion of "convergence"—a way to say that a sequence of vectors $\{s_n\}$ is getting closer and closer to some fixed vector s . The most natural way to interpret "closeness" of two vectors is to require that the distance between them is small, but to do this we must have a notion of distance in our vector space. As we have studied in class, a notion of distance can be defined for any vector space on which we have an inner product. Thus a natural place to

try to make sense of infinite sums of vectors is in an inner product space.¹

Let V be an inner product space. Can we find a set of vectors $\{v_n\}$, for $n = 1, 2, 3, \dots$ such that every vector v in V can be expressed as an infinite sum $\sum_{i=0}^{\infty} \alpha_i v_i$ in a unique way? Better yet, can we find an orthonormal set $\{v_n\}$ with this property? If so, such an orthonormal set is called an *orthonormal basis* for the space V , although we must be careful not to confuse this concept with a Hamel basis (in which we do not allow infinite sums). Fortunately, the answer is YES in many important applications. We now discuss a couple of these.

Fourier Series. One of the most important examples where this sort of alternative to the usual notion of a basis appears is in *Fourier Analysis*. In Fourier analysis, we are interested finding an “orthonormal basis” for the space of continuous functions on some interval.

Let $\mathcal{C}([-\pi, \pi])$ be the vector space of all continuous real valued functions on the interval $[-\pi, \pi]$. As we have seen, we can define an inner product on this space by

$$\langle f, g \rangle = \frac{1}{\pi} \int_{-\pi}^{\pi} fg \, dx.$$

Thus the distance between two functions f and g can be defined as

$$d(f, g) = \|f - g\| = \langle f - g, f - g \rangle^{1/2} = \left(\frac{1}{\pi} \int_{-\pi}^{\pi} (f - g)^2 dx \right)^{1/2}.$$

Now it makes sense to say that a sequence of continuous functions $\{s_n\}$ approaches some fixed function s . We mean that the distances $d(s, s_n)$ approach zero as n approaches infinity, that is, that the sequence of real numbers

$$\left(\frac{1}{\pi} \int_{-\pi}^{\pi} (s - s_n)^2 dx \right)^{1/2}$$

approaches zero as $n \rightarrow \infty$. (A subtle point to beware: this is not the same as saying that for each x , the sequence of real numbers $s_n(x)$ approaches the real number $s(x)$. The type of convergence we are interested in here is called “ L^2 convergence” or “convergence in the mean”, and should not be confused with this “pointwise convergence.”)

Given an expression $\sum_{i=1}^{\infty} f_n$ consisting of an infinite sum of functions, how can we tell if it is meaningful? We consider the sequence of partial sums $s_n = \sum_{i=1}^n f_n$, a sequence of functions in $\mathcal{C}([-\pi, \pi])$. The expression $\sum_{i=1}^{\infty} f_n$ is meaningful if the sequence $\{s_n\}$ converges (in the sense described above) to some function s in $\mathcal{C}([-\pi, \pi])$. In this case, we say that the infinite sum $\sum_{i=1}^{\infty} f_n$ exists and is equal to the function s .

¹Actually, there are vector spaces, called normed spaces, in which one can talk about distance, but which do not necessarily have an inner product. An inner product, after all, carries information not just about distance, but also about angles. Banach spaces, if you have heard of them, are important examples of normed spaces.

We now ask whether we can find an orthonormal set \mathcal{S} of functions $\{f_i\}$ for $i = 1, 2, 3, \dots$ such that every continuous function (ie, every vector in $\mathcal{C}([-\pi, \pi])$) is an infinite sum of the form $\sum_{i=1}^{\infty} \alpha_i f_i$ in a unique way. As we have proved in class, any orthonormal set in any inner product space will be linearly independent, and so such a set will be an orthonormal basis (but not necessarily a Hamel basis²). The answer is YES, and the set \mathcal{S} can be taken to be the set consisting of the functions

$$\begin{aligned} h_n(x) &= \sin(nx) \text{ for } n = 1, 2, 3, \dots \\ g_n(x) &= \cos(nx) \text{ for } n = 1, 2, 3, \dots \\ g_0(x) &= 1/2. \end{aligned}$$

It is easy to verify that these functions are orthonormal (by performing the necessary integrals). This gives us an orthonormal set of vectors in $\mathcal{C}([-\pi, \pi])$. What is also true, but not so obvious, is that *every continuous function on the interval $[-\pi, \pi]$ has a unique representation as an infinite linear combination of these functions*. This expression is called the *Fourier series* of a function and the coefficients of f_n that appear are called the *Fourier coefficients*. Interestingly, even though the space of all continuous functions on $[-\pi, \pi]$ has a countable orthonormal basis, it can be shown that it does not have a countable Hamel basis.

It turns out that for the space $\mathcal{C}([-\pi, \pi])$ defined above, the set \mathcal{S} is a maximal orthonormal set. In fact, an *orthonormal basis* for a infinite dimensional inner product space V can be defined as a *maximal orthonormal* set of vectors in V . Of course, as we have proven, the same definition is valid for finite dimensional inner product spaces as well. Essentially the same proof (using Zorn's Lemma) we used to show that every space has a Hamel basis shows that every inner product space V has an orthonormal basis. However, a drawback of this definition is that it is not clear that every vector in V has a unique expression as an infinite linear combination of the basis elements. In fact, this is not true in general. Please consult a book on functional analysis, such as [C], for more information.

One important difference between this type of orthonormal basis and a Hamel basis is that it is not true that *every* infinite linear combination of the basis elements belongs to our vector space. Indeed, an arbitrary infinite sum need not even converge to an actual function in general.

To learn more about Fourier Series, take a class on Fourier analysis such as Math 454 or see a book such as [BD].

Hilbert Spaces.

A Hilbert space is simply an inner product space with one additional property called completeness. Roughly speaking, the space is complete if every sequence of vectors which appears to be converging actually does converge. More precisely, the space is complete if every *Cauchy sequence* converges, where a sequence $\{s_n\}$ is Cauchy if for all positive numbers ϵ , there exists an integer N such that whenever $n, m > N$, the distance between s_n and s_m is less than ϵ .

²In fact, it can be shown that it is never a Hamel basis.

In general, just because a sequence appears to converge (that is, is Cauchy), we can not conclude that it actually does converge. For example, consider the vector space of continuous functions on the interval $[-\pi, \pi]$ with the inner product $\langle f, g, \rangle = \frac{1}{\pi} \int_{-\pi}^{\pi} fg dx$. The functions

$$\begin{aligned} f_n(x) &= 0 \text{ if } x \leq 0 \\ f_n(x) &= x^{1/n} \text{ if } x > 0 \end{aligned}$$

for $n = 1, 2, 3, \dots$ form a Cauchy sequence. This sequence appears to be converging, but it does *not* converge to a vector in $\mathcal{C}([-\pi, \pi])$. Indeed, this sequence converges to

$$\begin{aligned} f(x) &= 0 \text{ for } x \leq 0 \\ f(x) &= 1 \text{ for } x > 0, \end{aligned}$$

which is not continuous. Thus the sequence $\{f_n\}$ of vectors in $\mathcal{C}([-\pi, \pi])$ does not converge to a vector in $\mathcal{C}([-\pi, \pi])$. In this sense, the space $\mathcal{C}([-\pi, \pi])$ seems to be missing some elements—it is not complete.

It turns out that there is a way to define the “completion” of the space $\mathcal{C}([-\pi, \pi])$. This will be an inner product space which contains $\mathcal{C}([-\pi, \pi])$ and also contains every function which is a limit of continuous functions in the sense we just described. This completion space is a complete inner product space called $L^2([-\pi, \pi])$. The space L^2 turns out to consist of all “square integrable functions”—that is, all functions f defined on the interval $[-\pi, \pi]$ such that the integral $\int_{-\pi}^{\pi} f^2 dx$ is finite. The space $L^2([-\pi, \pi])$ is a prototypical example of a Hilbert space.

It turns out that the same trigonometric functions $h_n(x) = \sin(nx)$, $g_n(x) = \cos(nx)$, and $g_0(x) = 1/2$ form an orthonormal basis for $L^2([-\pi, \pi])$. This makes sense, since we already know that every continuous function is a limit of a sequence of functions obtained as finite linear combinations of these trigonometric functions, and we also know that every vector in L^2 is a limit of continuous functions. This emphasizes the difference between this kind of basis and a Hamel basis for an infinite dimensional space: two different spaces here have the same basis! The reason this can happen is that we do not allow *arbitrary* infinite linear combinations of the elements, but only those particular infinite sums that converge to a vector in the space in question.

Another important example of a Hilbert space is the space ℓ^2 of square-summable sequences as defined in Example 4 above. This is an inner product space, where the inner product is defined as a sort of “infinite dot product”:

$$\langle (\alpha_1, \alpha_2, \alpha_3, \dots), (\beta_1, \beta_2, \beta_3, \dots) \rangle = \sum_{i=1}^{\infty} \alpha_i \beta_i.$$

It is not obvious that this makes sense: why should such a sum necessarily converge? Of course it doesn’t converge in general if we take $(\alpha_1, \alpha_2, \alpha_3, \dots)$ and $(\beta_1, \beta_2, \beta_3, \dots)$ to be arbitrary elements in \mathbb{R}^{∞} , such as $(1, 1, 1, \dots)$ and $(2, 2, 2, \dots)$.

Fortunately, if we restrict to elements of ℓ^2 , the convergence can be proved using the Cauchy-Schwartz inequality. To see this, we need to show that the sequence of partial sums

$$s_n = \sum_{i=1}^n \alpha_i \beta_i$$

converges to some real number. It suffices to show the sum is absolutely convergent, so there is no harm in assuming all α_i and β_i are non-negative. By Cauchy-Schwartz for the standard inner product on \mathbb{R}^n , we see that

$$s_n \leq \|(\alpha_1, \dots, \alpha_n)\| \cdot \|(\beta_1, \dots, \beta_n)\|.$$

But as n goes to infinity, the quantity on the right goes to the product of the infinite sums $\sum_{i=1}^{\infty} \alpha_i^2$ and $\sum_{i=1}^{\infty} \beta_i^2$, which we are assuming are finite. Thus the sequence s_n is a bounded increasing sequence of real numbers, so it must have a limit. This shows that the sum $\sum_{i=1}^{\infty} \alpha_i \beta_i$ converges, so $\langle (\alpha_1, \alpha_2, \alpha_3, \dots), (\beta_1, \beta_2, \beta_3, \dots) \rangle$ is well-defined and ℓ^2 is an inner product space.

It turns out this inner product space is complete, so ℓ^2 is a Hilbert space. A good exercise is to show that the vectors $\{e_i\}$ (where e_i is the infinite sequence of zeros with a single one in the i -th spot) is an orthonormal basis for ℓ^2 .

There are many more interesting things that can be said about infinite dimensional inner product spaces. Please take a course in advanced analysis, such as Math 597 or Math 602, to learn more! You can also consult any book on functional analysis, such as [C] below.

REFERENCES

- [BD] Boyce, W. E. and DiPrima, *Elementary differential equations and boundary value problems*, John Wiley, New York, 1969.
- [HS] Hewitt, E., and Stromberg, K., *Real and Abstract Analysis*, Springer Verlag, New York, 1965.
- [C] Conway, John B., *A course in functional analysis*, Springer-Verlag, New York, 1990.