

Part III

The Computability Theory of Eigenvalues and Eigenvectors

Chapter 4

The Second Main Theorem, the Eigenvector Theorem, and Related Results

Introduction

In this chapter, we shall be mainly concerned with operators on Hilbert space, and especially with self-adjoint operators. We will assume that our operators are, in some natural sense, “effectively determined”. (The precise definition of this term is given in Section 1.) All of the standard operators of analysis and physics are effectively determined. However, we should emphasize that when an operator is called “effectively determined”, this designation applies only to the operator itself, and not to the quantities derived from it, such as its eigenvalues, eigenvectors, or spectrum.

We shall address the question: Which of the quantities associated with an “effectively determined” operator are computable? For example, are the eigenvalues computable? What about the sequence of eigenvalues? We shall see that the answer is “yes” for individual eigenvalues but “no” for the sequence of eigenvalues. More precisely, these statements hold for (bounded or unbounded) self-adjoint operators and for bounded normal operators.

Here we recall some distinctions set down in Chapter 0. When we assert, as we have done for self-adjoint operators, that the individual eigenvalues are computable but the sequence of eigenvalues need not be, we mean the following. For any fixed eigenvalue, we can program a computer to compute it. However, we might need a different program for each eigenvalue—i.e. there may be no master program which gives, for each n , the n -th eigenvalue.

One might ask whether the above results for self-adjoint operators can be extended to the case of a bounded linear operator on a Banach space. The answer is “no”, even for non-normal operators on Hilbert space. There exists an effectively determined bounded non-normal operator on Hilbert space which has a non-computable eigenvalue.

Let us return to self-adjoint operators, for which, as we have seen, the individual eigenvalues are computable. We can ask the same question for the eigenvectors. The answer is quite different. Even for an effectively determined *compact* self-adjoint operator, the eigenvectors associated with the eigenvalue $\lambda = 0$ need not be computable.

Here we should dispose of a triviality. Of course, for any operator, *some* of the eigenvectors are noncomputable—as we can see simply by taking a computable eigenvector and multiplying it by a noncomputable real. When we say that the eigenvectors associated with $\lambda = 0$ are not computable, we mean that *none* of them are computable.

This contrast between the computability of eigenvalues and eigenvectors can be given a physical interpretation. In quantum mechanics the eigenvalues are closely related to quantities actually measured—e.g. to the lines in the spectrum. By contrast, the eigenvectors are associated with the underlying state of the system. Our results show that the eigenvalues are computable, whereas the eigenvectors need not be.

Besides the eigenvalues and eigenvectors, we can ask similar questions about the spectrum. (For the definition of “eigenvalue” and “spectrum” see Section 1.) Again we are mainly concerned with the self-adjoint case. We shall show that there exists a computable sequence of real numbers which belong to the spectrum and whose closure coincides with the spectrum.

We now give a brief account of the sections in this chapter.

Section 1 contains the precise definition of “effectively determined” operator. It also gives a brief review of the notions of “self-adjointness”, “spectrum” and “eigenvalues” for bounded and unbounded operators.

Section 2 contains the Second Main Theorem, together with an investigation of several of its corollaries. The Second Main Theorem incorporates all of the results mentioned above for the spectra and eigenvalues of effectively determined (bounded or unbounded) self-adjoint operators. This theorem is best-possible, as we will eventually prove by suitable examples (Chapter 5, Section 8).

The proof of the Second Main Theorem is long and complicated, and it is deferred until Chapter 5.

Section 3 deals with discontinuities in the behavior of eigenvalues. For example, arbitrarily small perturbations of a self-adjoint operator can cause eigenvalues to abruptly disappear, while other eigenvalues—in quite different locations—are being suddenly created. Such discontinuities are frequently correlated with noncomputability. However, that is not the case here. Thus Section 3 (discontinuities in the eigenvalues) provides a counterpoint to Section 2 (computability of the eigenvalues).

Section 4 gives the example, promised above, of an effectively determined bounded non-normal operator with a noncomputable eigenvalue.

Sections 5 and 6 give the Eigenvector Theorem, together with its proof. This theorem asserts, as mentioned above, that there exists an effectively determined compact self-adjoint operator such that none of the eigenvectors corresponding to $\lambda = 0$ are computable. The proof of the Eigenvector Theorem is somewhat indirect. It begins, in Section 5, with a construction based on an ad-hoc (i.e. “artificial”) computability structure. (Cf. Chapter 2, Section 7.) Then in Section 6 we show how to translate this ad-hoc construction into one involving the “natural” intrinsic computability structure of $L^2[0, 1]$.

Section 7 ties up some loose ends, and also, for the first time in this chapter, deals with Banach spaces other than Hilbert space. This section contains two main results. First it gives a proof of the Effective Independence Lemma, which asserts that, from

any effective generating set $\{e_n\}$, we can extract a linearly independent effective generating subset. This lemma plays a role in our proof of the Eigenvector Theorem. While the result is hardly surprising, its proof is not quite so easy as might be supposed.

Second, we address the question: Are all effectively separable computability structures on a Banach space X related via (not necessarily computable) isometries? The answer turns out to be “yes” for Hilbert space, but “no” for Banach spaces in general. The “yes” part is also a step on the way to the Eigenvector Theorem (Lemma 8, below), while the “no” part is established by a counterexample given at the end of Section 7.

1. Basic Notions for Unbounded Operators, Effectively Determined Operators

Throughout most of this chapter, the underlying Banach space will be an effectively separable Hilbert space H . The inner product of two vectors $x, y \in H$ is denoted by (x, y) .

We recall from Chapter 3 the notion of a closed unbounded operator. An operator $T: H \rightarrow H$ is called *closed* if T has a closed graph. In general, the domain of T is not H , but a dense linear subspace $\mathcal{D}(T)$ of H . Thus in order to define T we must first specify the domain $\mathcal{D}(T)$ and then describe the action of T on this domain.

We now define the adjoint T^* of T . In order to motivate what follows, it is useful to begin with the familiar case of bounded operators. Let $T: H \rightarrow H$ be a bounded linear operator on H . Then, as is well known, the *adjoint* operator T^* is defined by

$$(Tx, y) = (x, T^*y) \quad \text{for all } x, y \in H.$$

The definition of the adjoint for unbounded T is a natural extension of this. We must define the domain $\mathcal{D}(T^*)$ and the action of T^* on this domain.

Definition (first variant). Let $T: H \rightarrow H$ be a closed operator with domain $\mathcal{D}(T)$.

a) A vector y belongs to the domain $\mathcal{D}(T^*)$ of T^* if there exists a vector z such that:

$$(Tx, y) = (x, z) \quad \text{for all } x \in \mathcal{D}(T).$$

b) When such a z exists, we define T^*y to be z .

Note. Thus we have the identity $(Tx, y) = (x, T^*y)$, just as in the bounded case.

It is well known (see e.g. Riesz, Sz.-Nagy [1955]) that T^* is well-defined, closed, and that its domain $\mathcal{D}(T^*)$ is dense in H .

We have preferred the above definition of T^* because it shows clearly the connection with the familiar bounded case. However, for serious work, an equivalent definition based on graphs turns out to be more powerful. First we recall that

the graph of T consists of all pairs $\langle x, Tx \rangle$ with $x \in \mathcal{D}(T)$. [Here we write $\langle \ , \ \rangle$ for ordered pairs to avoid confusion with the inner product $(\ , \)$.] Now the second equivalent definition is:

Definition (second variant). Let $T: H \rightarrow H$ be a closed operator, and let G be the graph of T in $H \times H$. Let G^\perp be the orthogonal complement of G in $H \times H$. Then:

a) A vector y belongs to the domain $\mathcal{D}(T^*)$ if there exists a vector z such that the pair $\langle z, y \rangle \in G^\perp$.

b) In this case, we define T^*y to be $-z$.

The proof of equivalence of the two preceding definitions is a routine exercise (cf. Riesz, Sz.-Nagy [1955]). The “ $-$ ” in part (b) of the second definition is not a misprint. It occurs because of formal manipulations involving the inner product on $H \times H$.

We now come to the basic:

Definition. A closed operator $T: H \rightarrow H$ (bounded or unbounded) is said to be *self-adjoint* if it coincides with its adjoint, i.e. if $T = T^*$.

In the unbounded case, it is important to stress that two operators are considered to be “equal” if and only if they have the same graph. As we recall from Chapter 3, an unbounded closed operator T_1 may possess a proper extension T_2 : this means that the domain $\mathcal{D}(T_1) \subsetneq \mathcal{D}(T_2)$, although T_2 coincides with T_1 on their common domain of definition.

The most important operators of analysis and physics—and, in particular, the so-called “observables” of quantum mechanics—are either self-adjoint or possess self-adjoint extensions. Throughout the remainder of this book, unless stated otherwise, we shall be concerned with (bounded or unbounded) operators which are self-adjoint.

There is one minor exception. In the bounded case, we shall sometimes consider normal operators. A bounded operator $T: H \rightarrow H$ is said to be *normal* if it commutes with its adjoint, i.e. if $T^*T = TT^*$. Normal operators possess many of the same properties as self-adjoint operators. In particular, the Second Main Theorem also holds for bounded normal operators, as we shall prove in Section 6 of Chapter 5.

Eigenvalues and spectrum

The structure of a bounded or unbounded self-adjoint operator is determined to a substantial degree by the eigenvalues and spectrum of that operator.

Definition (spectrum). Let T be a closed operator (bounded or unbounded). A number λ belongs to the *spectrum* of T if the operator $(T - \lambda)$ does not have a bounded inverse.

We observe that, for λ not to be in spectrum (T) , the inverse $(T - \lambda)^{-1}$ must be bounded, whether T itself is bounded or not.

It is well known (cf. Riesz, Sz.-Nagy [1955]) that the spectrum is a closed set, and that when T is self-adjoint, the spectrum is real. Hence, for self-adjoint T , the inverse $(T - \lambda)^{-1}$ exists and is bounded for all $\lambda = \alpha + i\beta$ with $\beta \neq 0$; in particular, $(T - i)^{-1}$ exists. In addition, for self-adjoint T , the spectrum is a bounded set if and only if T is bounded.

Definition (eigenvalues). A number λ is called an *eigenvalue* of T if there exists a nonzero vector x (the corresponding *eigenvector*) such that $Tx = \lambda x$.

The eigenvalues form a subset of the spectrum, which in general is a proper subset. Points of the spectrum which are not eigenvalues are commonly said to belong to “the continuous spectrum”.

The eigenvalues and continuous spectrum have a direct physical significance in quantum mechanics. For example, when the emissions from a light source are observed in the spectroscope, the eigenvalues are closely related to the appearance of “bright lines”, whereas the continuous spectrum is reflected in the presence of continuous bands of light.

Effectively determined operators

Up to now, there has been no mention within this section of computability. We have simply set down some standard facts from functional analysis. Now we must define what it means for an operator to “act effectively”. We shall call such operators *effectively determined*.

Recall that, by assumption, the space H is effectively separable. This means that there is a computability structure defined on H , and that H has an effective generating set $\{e_n\}$ —i.e. a computable sequence $\{e_n\}$ whose linear span is dense in H . (See Chapter 2.)

Operators map vectors into vectors. Thus to “determine” an operator means to know, in some effective manner, how the operator acts on computable vectors and computable sequences of vectors. In the unbounded case, we need to know one thing more. We have seen that, to define an unbounded operator, we really need to know its graph. Thus we should expect the definition of an “effectively determined” unbounded operator to involve an effective specification of the graph.

Since the graph is contained within $H \times H$, we must first define “computability” for the cartesian product. The definition is obvious. A sequence of points $\langle x_n, y_n \rangle \in H \times H$ is called *computable* if $\{x_n\}$ and $\{y_n\}$ are computable in H .

Definition (Effectively determined operator). A closed operator $T: H \rightarrow H$ is *effectively determined* if there is a computable sequence $\{e_n\}$ in H such that the pairs $\{\langle e_n, Te_n \rangle\}$ form an effective generating set for the graph of T .

We recall that, by the definition of “effective generating set”, this means that $\{\langle e_n, Te_n \rangle\}$ is computable in $H \times H$, and that the linear span of $\{\langle e_n, Te_n \rangle\}$ is dense in the graph of T . In particular, this implies that $\{e_n\}$ and $\{Te_n\}$ are computable and that the span of $\{e_n\}$ is dense in H . Thus, as a corollary of the above definition,

$\{e_n\}$ is also an effective generating set for H . The converse of this corollary is false, however. When the operator T is unbounded, the density of (linear span of $\{e_n\}$) in H does *not* imply that $\{\langle e_n, Te_n \rangle\}$ spans a dense subset of the graph.

The following may help to explain what could go wrong. We saw in Chapter 3, Section 1 that a closed operator T_1 can have a proper closed extension T_2 , where $\mathcal{D}(T_1) \subsetneq \mathcal{D}(T_2)$. Suppose we are really interested in the extended operator T_2 . Suppose, however, that we choose an effective generating set $\{e_n\}$ contained within the smaller domain $\mathcal{D}(T_1)$. Now $\mathcal{D}(T_1)$ is dense in $\mathcal{D}(T_2)$, since it is dense in H . Thus $\{e_n\}$ is a perfectly good effective generating set for H . But on $H \times H$, the pairs $\langle e_n, Te_n \rangle$ span at most the graph of T_1 . For they lie within the graph of T_1 , and this graph is *closed*. Thus $\{\langle e_n, Te_n \rangle\}$ cannot approach those points on graph (T_2) which lie outside of graph (T_1).

Of course, these difficulties can only occur for unbounded operators. Consider, by contrast, a bounded operator T . Then T is continuous, and the convergence in H of $\{e_n\}$ (or linear combinations thereof) automatically implies convergence of the pairs $\{\langle e_n, Te_n \rangle\}$.

In summary, a bounded operator T is effectively determined if and only if it maps an effective generating set $\{e_n\}$ for H onto a computable sequence $\{Te_n\}$.

As mentioned above, all of the standard operators of analysis and physics are effectively determined. The most interesting of these are unbounded.

2. The Second Main Theorem and Some of Its Corollaries

We recall that H denotes an effectively separable Hilbert space—i.e. a Hilbert space with a computability structure for which there is an effective generating set $\{e_n\}$.

Second Main Theorem. *Let $T: H \rightarrow H$ be an effectively determined (bounded or unbounded) self-adjoint operator. Then there exists a computable sequence of real numbers $\{\lambda_n\}$ and a recursively enumerable set A of natural numbers such that:*

- i) *Each $\lambda_n \in \text{spectrum}(T)$, and the spectrum of T coincides with the closure of $\{\lambda_n\}$.*
- ii) *The set of eigenvalues of T coincides with the set $\{\lambda_n: n \in \mathbb{N} - A\}$. In particular, each eigenvalue of T is computable.*
- iii) *Conversely every set which is the closure of $\{\lambda_n\}$ as in (i) above occurs as the spectrum of an effectively determined self-adjoint operator.*
- iv) *Likewise, every set $\{\lambda_n: n \in \mathbb{N} - A\}$ as in (ii) above occurs as the set of eigenvalues of some effectively determined self-adjoint operator T . If the set $\{\lambda_n\}$ is bounded, then T can be chosen to be bounded.*

Note. Concerning the boundedness of T : In (iii), where $\{\lambda_n\}$ determines the entire spectrum, the boundedness of $\{\lambda_n\}$ implies the boundedness of T —for, as is well known (Riesz, Sz.-Nagy [1955]), the spectral norm of a self-adjoint operator coincides with its norm. However, in (iv), where $\{\lambda_n\}$ gives only the set of eigenvalues (and not the entire spectrum), the boundedness of T must be considered separately.

The proof of this theorem is long and complicated. It is deferred until Chapter 5. In fact, it forms the entire content of Chapter 5. We turn now to some consequences of the Second Main Theorem.

As stated in (ii) above, the individual eigenvalues of T are computable. However, the sequence of eigenvalues need not be, as we now show.

Theorem 1 (The sequence of eigenvalues). *There exists an effectively determined bounded self-adjoint operator $T: H \rightarrow H$ whose sequence of eigenvalues is not computable.*

Proof. We use the counterexample asserted in (iv) above. (For its details, see Chapter 5, Section 8.) Let $\{\lambda_n\}$ be the computable sequence $\lambda_n = 2^{-n}$, and let A be any recursively enumerable non recursive set of integers. Then by (iv) there exists an effectively determined bounded self-adjoint operator T whose eigenvalues coincide with the set $\{\lambda_n: n \in \mathbb{N} - A\}$. But now any effective listing of these λ_n would give an effective listing of the corresponding set of n 's, i.e. an effective listing of the complement of A . Since A is not recursive, this is impossible. \square

For the special case of compact operators, the phenomenon exhibited in Theorem 1 cannot occur. Namely, as a consequence of (i) above, we have:

Theorem 2 (Compact operators). *Let $T: H \rightarrow H$ be an effectively determined compact self-adjoint operator. Then the set of eigenvalues of T forms a computable sequence of real numbers.*

Proof. It is well known (Riesz, Sz.-Nagy [1955]) that, if T is compact, the spectrum of T consists of isolated eigenvalues $\lambda \neq 0$ together with $\{0\}$ as their only possible limit point. Now take the sequence $\{\lambda_n\}$ given by (i). Since the eigenvalues $\lambda \neq 0$ are isolated, and $\{\lambda_n\}$ is dense in the spectrum, it follows that every eigenvalue $\lambda \neq 0$ equals λ_n for some n .

We now consider the value $\lambda = 0$, which may or may not be an eigenvalue. To deal with this, we first extract the computable subsequence $\{\lambda'_n\}$ of all $\lambda_n \neq 0$, and then include or exclude the value $\lambda = 0$ according as it is an eigenvalue or not. \square

Just as the negative result (iv) gave information about the sequence of eigenvalues, the negative result (iii) gives information about the operator norm.

Theorem 3 (The operator norm). *There exists an effectively determined bounded self-adjoint operator $T: H \rightarrow H$ whose norm is not a computable real.*

Proof. We recall that the “norm” of an operator T is defined to be $\sup \{\|Tx\|/\|x\|: x \neq 0\}$. We shall also use the “spectral norm” defined as $\sup \{|\lambda|: \lambda \in \text{spectrum}(T)\}$. It is well known (Riesz, Sz.-Nagy [1955], Halmos [1951]) that for self-adjoint operators, the norm and the spectral norm coincide.

Now let $a: \mathbb{N} \rightarrow \mathbb{N}$ be a one to one recursive function listing a recursively enumerable non recursive set A . Let

$$\lambda_n = \sum_{k=0}^n 2^{-a(k)}.$$

Then $\{\lambda_n\}$ is a computable monotone sequence converging noneffectively to a non-computable real number α (cf. Chapter 0, Sections 1 and 2).

Let T be the operator, promised in (iii) above, whose spectrum is the closure of $\{\lambda_n\}$. Then the spectral norm of T (= the norm of T) is $\sup \{\lambda_n\} = \alpha$, a noncomputable real. \square

3. Creation and Destruction of Eigenvalues

It is a well-known fact that eigenvalues of a self-adjoint operator can be instantaneously created and destroyed—that is, their behavior can be highly discontinuous. Such discontinuities often lead to noncomputability. Eigenvalues furnish an exception to this rule. For, as this section will show, we have discontinuity. By contrast, the key theorem of the last section asserts that the eigenvalues are computable.

What do we mean here by discontinuity? We mean that there is a one-parameter family T_ε of bounded self-adjoint operators, such that T_ε varies continuously with ε , but the behavior of the eigenvalues is discontinuous.

What do we mean when we say that T_ε “varies continuously with ε ”? We mean that we have continuity in the sense of uniform convergence on the unit ball, i.e. in the operator norm defined by $\|T\| = \sup \{\|Tx\|: \|x\| \leq 1\}$. Thus to say that $T_\varepsilon \rightarrow T$ as $\varepsilon \rightarrow 0$ means that the operator norm $\|T_\varepsilon - T\| \rightarrow 0$. This is a very strong condition: in fact the operator norm topology is the strongest of the topologies commonly employed for operators.

Theorem 4 (Creation and destruction of eigenvalues). *There exists a one parameter family $\{T_\varepsilon\}$ of bounded self-adjoint operators on a separable Hilbert space H such that:*

- i) T_ε varies continuously with ε in terms of the operator norm topology on $\{T_\varepsilon\}$.
- ii) For $\varepsilon = 0$, the operator T_0 has the unique eigenvalue $\lambda = 0$, and this eigenvalue is of multiplicity one.
- iii) For all $\varepsilon \neq 0$ and sufficiently close to zero, (a) T_ε has no eigenvalue near zero, (b) T_ε has an eigenvalue near each of the points ± 1 , and (c) all eigenvalues are of multiplicity one.

Proof. Let H be the direct sum of $L^2[-1, 1]$ and an element δ of norm one generating a 1-dimensional Hilbert space (δ). Thus δ is orthogonal to $L^2[-1, 1]$. We denote functions in $L^2[-1, 1]$ by the letters f, g, h, \dots . We define the operator

T_ε on H by:

$$T_\varepsilon[f(x)] = xf(x) + \varepsilon \cdot \int_{-1}^1 f(x) dx \cdot \delta,$$

$$T_\varepsilon[\delta] = \varepsilon \cdot 1 \quad (\text{a constant function on } [-1, 1]).$$

It is easy to verify that T_ε is self-adjoint. Also, we see at once that T_ε varies continuously with ε in terms of the operator norm.

Now for $\varepsilon = 0$, T has the eigenvalue 0 with eigenvector δ .

We shall show that for all sufficiently small $\varepsilon \neq 0$, T_ε has an eigenvalue of multiplicity one located near each of the points $\lambda = \pm 1$, and no eigenvalue near 0. Thus the eigenvalue 0 is destroyed, whereas λ near ± 1 are created.

Take $\varepsilon \neq 0$. Now any eigenvector of T_ε must involve δ . Hence, multiplying by a scalar, we may assume that the eigenvector has the form $f(x) + \delta$. Let λ be the corresponding eigenvalue, so that

$$T_\varepsilon[f(x) + \delta] = \lambda[f(x) + \delta].$$

Recalling the definition of T_ε , and equating the $L^2[-1, 1]$ and (δ) components, we obtain:

$$xf(x) + \varepsilon = \lambda f(x),$$

$$\varepsilon \cdot \int_{-1}^1 f(x) dx = \lambda.$$

But the first equation gives

$$f(x) = -\varepsilon/(x - \lambda),$$

and from the second equation we have

$$-\varepsilon^2 \int_{-1}^1 \frac{dx}{x - \lambda} = \lambda.$$

We now focus our attention on the last two equations. We must determine the values of λ which satisfy these equations, and for which the associated $f(x) \in L^2[-1, 1]$. All values $\lambda \in [-1, 1]$ are ruled out, since the function $f(x) = -\varepsilon/(x - \lambda)$ is not L^2 on $[-1, 1]$. Thus we need examine only $\lambda > 1$ and $\lambda < -1$. Let us consider $\lambda > 1$; the other case is similar. For $\lambda > 1$, the function $f(x)$ is L^2 on $[-1, 1]$, and so we need only ask whether or not the above equations are satisfied. We claim that, for each $\varepsilon \neq 0$, there is a unique $\lambda > 1$ which satisfies the last displayed equation above. Furthermore, $\lambda \downarrow 1$ as $\varepsilon \rightarrow 0$. To see this, we rewrite this equation as

$$\int_{-1}^1 \frac{dx}{x - \lambda} = \frac{-\lambda}{\varepsilon^2}.$$

Fix any $\varepsilon \neq 0$. Then as $\lambda \downarrow 1$, the integral above decreases to $-\infty$, whereas the fraction increases to $-1/\varepsilon^2$. Hence, by the intermediate value theorem, there is a unique solution λ . \square

4. A Non-normal Operator with a Noncomputable Eigenvalue

In the Second Main Theorem and its extensions we assert that the eigenvalues of an effectively determined self-adjoint (or normal) operator are computable. Since the proof (cf. Chapter 5) uses the spectral theorem, the result is closely tied to the assumed self-adjointness/normality of the operator. This restriction is necessary as we now show.

The following theorem is proved for the Hilbert space $H = L^2[0, 1]$ with its standard computability structure. However, as we will show later (Lemma 8, Section 6), the result could be transferred to any effectively separable Hilbert space.

Theorem 5 (Noncomputable eigenvalues). *There exists an effectively determined bounded operator $T: H \rightarrow H$ (not self-adjoint or normal) which has a noncomputable real number α as an eigenvalue.*

Proof. Let $a: \mathbb{N} \rightarrow \mathbb{N}$ be a one to one recursive function which enumerates a recursively enumerable non recursive set A . Let $H = L^2[0, 1]$, and let $\{e_n\}$ be a computable orthonormal basis for H .

To define T , it suffices to give the value of $T(e_n)$ for all n . We define, for each n :

$$T(e_n) = \left(\sum_{k=0}^n 4^{-a(k)} \right) e_n + 2^{-a(n)} \cdot \sum_{k=0}^{n-1} 2^{-a(k)} e_k.$$

Now we show that T is bounded. Write $T = T_1 + T_2$, where T_1 and T_2 correspond respectively to the first and second terms in the above expression for $T(e_n)$. Then T_1 is a bounded self-adjoint operator: in terms of the basis $\{e_n\}$, it corresponds to a diagonal matrix with the bounded sequence of eigenvalues $\left\{ \sum_0^n 4^{-a(k)} \right\}$. For T_2 we reason as follows. The vectors $\sum_0^{n-1} 2^{-a(k)} e_k$ are bounded in norm by $\sum_0^n 2^{-a(k)} \leq 2$. Hence for each n , $\|T_2(e_n)\| \leq 2 \cdot 2^{-a(n)}$. Take an arbitrary vector x :

$$x = \sum_{n=0}^{\infty} c_n e_n.$$

Since $\{e_n\}$ is an orthonormal basis, the norm $\|x\|$ is just the l^2 -norm of the sequence $\{c_n\}$. Now

$$\|T_2 x\| \leq \sum_{n=0}^{\infty} |c_n| \|T_2 e_n\| \leq 2 \sum_{n=0}^{\infty} |c_n| \cdot 2^{-a(n)}.$$

Since $\{2^{-a(n)}\}$ is an l^2 -sequence, it follows from the Schwarz inequality (for sequences) that T_2 is bounded. Hence $T = T_1 + T_2$ is bounded.

Now we show that T has a noncomputable eigenvalue. Namely, we show that the vector

$$\sum_{k=0}^{\infty} 2^{-a(k)} e_k$$

is an eigenvector with the eigenvalue

$$\alpha = \sum_{k=0}^{\infty} 4^{-a(k)}.$$

First we note the identity:

$$T \left[\sum_{k=0}^n 2^{-a(k)} e_k \right] = \left(\sum_{k=0}^n 4^{-a(k)} \right) \cdot \left(\sum_{k=0}^n 2^{-a(k)} e_k \right).$$

[This identity follows by a straightforward induction on n , using the definition of $T(e_n)$ given above. Of course, this identity is also the motivation for our definition of $T(e_n)$.]

Letting $n \rightarrow \infty$, we deduce that

$$T \left[\sum_{k=0}^{\infty} 2^{-a(k)} e_k \right] = \left(\sum_{k=0}^{\infty} 4^{-a(k)} \right) \cdot \left(\sum_{k=0}^{\infty} 2^{-a(k)} e_k \right).$$

Finally, the eigenvalue $\alpha = \sum_0^{\infty} 4^{-a(k)}$ is not a computable real. \square

Note. These last few steps show why it is essential that the operator T be non-normal. For the above construction involves a sequence of eigenvectors $\left\{ \sum_{k=0}^n 2^{-a(k)} e_k \right\}$, all very close together and converging to $\sum_0^{\infty} 2^{-a(k)} e_k$. These eigenvectors have the slightly different eigenvalues $\sum_{k=0}^n 4^{-a(k)}$. With a normal operator, distinct eigenvalues would force the vectors to be orthogonal and not close together.

5. The Eigenvector Theorem

In this and the following section we will prove:

Theorem 6 (The Eigenvector Theorem). *Let $H = L^2[0, 1]$ with its intrinsic computability structure. There exists an effectively determined compact self-adjoint operator $T: H \rightarrow H$ with the following properties.*

(1) *The number $\lambda = 0$ is an eigenvalue of T of multiplicity one (i.e. the space of eigenvectors corresponding to $\lambda = 0$ is one dimensional).*

(2) *None of the eigenvectors corresponding to $\lambda = 0$ is computable.*

As noted in the Introduction to this chapter, the proof is indirect. In this section we shall prove the following weaker result.

Eigenvector Theorem (preliminary form). *There exists an ad hoc computability structure on H , and an operator $T: H \rightarrow H$ which is effectively determined in terms of this ad-hoc structure and such that T is compact, self-adjoint, and satisfies conditions (1) and (2) above. Furthermore, H is effectively separable in terms of this structure.*

In Section 6 we shall show how to translate this preliminary theorem into its desired final form (Theorem 6 above), involving the (“natural”) intrinsic computability structure on $L^2[0, 1]$.

Remarks. We recall (Chapter 2, Section 7) that an *ad hoc* computability structure is a non-intrinsic structure—i.e. a structure that can be regarded as “artificial”—which nevertheless satisfies the axioms for computability on a Banach space. The proof of the Eigenvector Theorem provides the main application of ad hoc computability given in this book. This approach is by no means an exercise in “fancy” technique. A bit of explanation seems in order.

In the Eigenvector Theorem, we have two objects to deal with: (a) the operator T , and (b) the computability structure. Obviously, to build such a counterexample, at least one of these must be somewhat intricate. In the preliminary (ad hoc) construction, the operator T is very simple, and it is the computability structure which carries the essential ideas of the counterexample. In Section 6 we reverse field, showing how *any* such counterexample can be translated into one involving the natural computability structure and a complicated operator.

Of course, the final form of the theorem (as completed in Section 6) is the result we mainly want, since it is the natural computability structure on $L^2[0, 1]$ which is of primary interest. However, in this final form, the operator T becomes so complicated that its intuitive meaning is lost. It seems that the final form of the operator would be much harder to discover, ab initio, than the slightly perturbed computability structure with which we begin this construction.

Proof of the Eigenvector Theorem (preliminary form). We begin by defining the ad hoc computability structure on $H = L^2[0, 1]$. First we take any one of the standard computable orthonormal bases $\{e_n\}$ for $L^2[0, 1]$: e.g. let $\{e_n\}$ be the sequence of functions $\{e^{2\pi imx}\}$, with the integers $m = 0, \pm 1, \pm 2, \dots$ mapped onto the natural numbers $n = 0, 1, 2, \dots$ in a standard computable way. For convenience, assume that $m = 0$ corresponds to $n = 0$.

Let H_0 be the closed subspace of H spanned by the vectors $\{e_1, e_2, \dots\}$; thus H_0 consists of all vectors in H which are orthogonal to e_0 . The vector e_0 will play a special role in our construction, and therefore for typographical clarity we write:

$$\Lambda = e_0.$$

[We emphasize that there is nothing “exotic” about Λ ; in fact $\Lambda = e_0 = e^{2\pi i 0x}$ is just the constant function 1.]

Now the computability structure which we shall place on H will coincide with the natural structure on the subspace H_0 , but will behave in an “artificial” manner with respect to the vector Λ . In this ad hoc structure, the vector Λ will not be computable. Here are the details.

Let $a: \mathbb{N} \rightarrow \mathbb{N}$ be a recursive function which enumerates a recursively enumerable non-recursive set A in a one to one manner. We assume that $0 \notin A$. We define a sequence of positive reals $\{\alpha_n\}$ and a single positive real number γ by setting:

$$\alpha_n = 2^{-a(n)} \quad \text{for } n \geq 1,$$

$$\gamma^2 = 1 - \sum_{n=1}^{\infty} \alpha_n^2.$$

It is important to observe that γ is not computable. For if it were, then the series $\sum \alpha_n^2$ would converge effectively, contradicting the fact that A is not recursive.

Now we define the vector

$$f = \gamma\Lambda + \sum_{n=1}^{\infty} \alpha_n e_n.$$

We observe that $\|f\| = 1$.

Remarks. In terms of the natural computability structure on $L^2[0, 1]$, the vector f is not computable (since the coefficient γ is not computable). In the ad hoc structure below, we will declare f to be “computable”. We will declare, further, that the standard notion of computability holds on H_0 . This essentially determines the ad hoc structure. It takes some work, however, to show that this structure satisfies the axioms for computability on a Banach space. (Cf. Lemmas 1, 2, and 3 below.)

We put an ad hoc computability structure on H by specifying that

$$\{f, e_1, e_2, e_3, \dots\}$$

is an effective generating set.

Equivalently, a sequence $\{x_n\}$ of vectors in H is *ad hoc computable* if:

$$x_n = \beta_n f + y_n,$$

where $\{\beta_n\}$ is a computable sequence of complex numbers, and $\{y_n\}$ is a computable sequence of vectors (in the standard sense) in H_0 .

As mentioned above, we will prove in Lemma 3 that this definition satisfies the axioms for a computability structure.

The operator T. We define T by setting:

$$T\Lambda = 0,$$

$$Te_n = 2^{-n}e_n \quad \text{for } n \geq 1.$$

Since $\{\Lambda = e_0, e_1, e_2, \dots\}$ form an orthonormal basis for H , T is self-adjoint. Since the sequence $2^{-n} \rightarrow 0$, T is compact. Clearly 0 is an eigenvalue of T of multiplicity one; the corresponding eigenvectors are the scalar multiples of Λ .

We shall show in Lemma 4 that T is effectively determined in terms of the ad hoc computability structure which we have defined on H .

We show in Lemma 5 that no nonzero multiple of the eigenvector Λ is computable in terms of the ad hoc structure.

By combining Lemmas 3, 4, and 5, we obtain all of the conditions stated in the preliminary form of the Eigenvector Theorem.

Statement and proof of Lemmas 1–5. We begin by noting that some of these lemmas are stated in the complete generality of an arbitrary Hilbert space H with a computability structure. In these cases, obviously, we must use nothing but the axioms for computability on a Banach space. When H is a specific space (e.g. $H = L^2[0, 1]$) we shall say so.

The following is important, but hardly deserves to be called a lemma. Let H be any Hilbert space with a computability structure. If $\{x_n\}$ and $\{y_n\}$ are computable sequences of vectors in H , then the double sequence of inner products $\{(x_n, y_m)\}$ is computable. For

$$(x_n, y_m) = \frac{1}{4}[\|x_n + y_m\|^2 - \|x_n - y_m\|^2 + i\|x_n + iy_m\|^2 - i\|x_n - iy_m\|^2].$$

The Linear Forms Axiom gives us the computability in H of the double sequences $\{x_n \pm y_m\}$ and $\{x_n \pm iy_m\}$. Then the Norm Axiom implies that $\{(x_n, y_m)\}$ is computable. (Of course, to handle the double sequences, we use one of the standard recursive pairing functions from $\mathbb{N} \times \mathbb{N}$ onto \mathbb{N} .)

Lemma 1. *Let H be a Hilbert space with an arbitrary computability structure imposed upon it. Suppose there exists a computable orthonormal basis $\{e_n\}$ for H . Take any sequence of vectors $\{x_n\}$, given by*

$$x_n = \sum_{k=0}^{\infty} c_{nk} e_k.$$

Then the sequence $\{x_n\}$ is computable in H if and only if:

- i) *the double sequence $\{c_{nk}\}$ of “Fourier coefficients” is computable, and*
- ii) *the series $\sum_{k=0}^{\infty} |c_{nk}|^2$ converges effectively in k and n .*

Proof. The “if” part is trivial: the computability of $\{x_n\}$ follows immediately from the Linear Forms and Limit Axioms. For the “only if”, take a computable sequence $\{x_n\}$. As we have seen, the sequence of inner products (x_n, e_k) is computable; this gives (i). To obtain (ii), we use the Norm Axiom. This tells us that $\|x_n\|^2 = \sum_{k=0}^{\infty} |c_{nk}|^2$ is a computable sequence of reals. Since the sequences of partial sums $\sum_{i=0}^k |c_{ni}|^2$ are

computable and converge monotonically to a computable sequence of reals, the convergence is effective in k and n . \square

The next lemma embodies one of the key steps on the way to proving that the ad hoc structure on $L^2[0, 1]$ satisfies the axioms. The point is that the vector x considered there is *not* assumed to be computable. In terms of Lemma 1 above, x satisfies condition (i) but not necessarily condition (ii).

Lemma 2. *Let H be any Hilbert space with a computability structure and with a computable orthonormal basis $\{e_n\}$ for H . Let $\{y_m\}$ be any computable sequence of vectors in H . Take any vector $x = \sum c_k e_k$ for which the sequence of “Fourier coefficients” $\{c_k\}$ is computable (condition (i) of Lemma 1). Then the sequence of inner products $\{(x, y_m)\}$ is computable.*

Proof. We have not assumed that the series $\sum |c_k|^2$ for $\|x\|^2$ converges effectively. However, there exists an integer M which bounds $\|x\|$.

Now let the “Fourier expansions” of the y_m be

$$y_m = \sum_{k=0}^{\infty} d_{mk} e_k.$$

Since $\{y_m\}$ is computable, the double sequence $\{d_{mk}\}$ is computable, and the series $\sum_{k=0}^{\infty} |d_{mk}|^2$ converge effectively in k and m by Lemma 1. Furthermore,

$$(x, y_m) = \sum_{k=0}^{\infty} c_k \bar{d}_{mk} \quad \text{for all } m.$$

Since the sequences $\{c_k\}$ and $\{d_{mk}\}$ are computable, to prove that $\{(x, y_m)\}$ is computable we have only to show that in the above series the convergence is effective in k and m .

Since $\sum_{k=0}^{\infty} |d_{mk}|^2$ converges effectively, there is a recursive function $e(m, N)$ such that

$$\left(\sum_{k=e(m, N)}^{\infty} |d_{mk}|^2 \right)^{1/2} \leq \frac{2^{-N}}{M} \quad \text{for all } m, N.$$

Since the norm of x is dominated by M , the Schwarz inequality implies that

$$\sum_{k=e(m, N)}^{\infty} |c_k \bar{d}_{mk}| \leq 2^{-N} \quad \text{for all } m, N.$$

Thus to compute $(x, y_m) = \sum_{k=0}^{\infty} c_k \bar{d}_{mk}$ to within an error of 2^{-N} , we merely compute the first $e(m, N)$ terms of the series. \square

From Lemma 2 we derive:

Lemma 3. *The “ad hoc” computability structure defined above on H satisfies the axioms for computability on a Banach space. Furthermore, this computability structure is effectively separable.*

Proof. Only the Norm Axiom causes any difficulty. Given a computable sequence $\{x_n\}$ in H , we need to compute $\{\|x_n\|\}$. By definition:

$$x_n = \beta_n f + y_n,$$

with $\{\beta_n\}$ computable in \mathbb{C} , and $\{y_n\}$ computable in H_0 . We recall that

$$f = \gamma \Lambda + \sum_{n=1}^{\infty} \alpha_n e_n,$$

where $\{\alpha_n\}$ is computable, but the series $\sum \alpha_n^2$ is not effectively convergent, and γ is a noncomputable real adjusted so that $\|f\| = 1$. Then:

$$\|x_n\|^2 = |\beta_n|^2 + \|y_n\|^2 + 2 \cdot \operatorname{Re} [\beta_n (f, y_n)].$$

We know that $\{\beta_n\}$, $\{|\beta_n|^2\}$, and $\{\|y_n\|^2\}$ are computable. Thus it suffices to show that

$$\{(f, y_n)\}$$

is computable. Let g be the projection of f on H_0 , namely:

$$g = \sum_{n=1}^{\infty} \alpha_n e_n.$$

Then, since $y_n \in H_0$, $(f, y_n) = (g, y_n)$ for all n . Now g is not computable in H_0 . However, the “Fourier coefficients” $\{\alpha_n\}$ of g are computable, and the sequence of vectors $\{y_n\}$ is computable. Hence, by Lemma 2, the sequence of inner products $\{(g, y_n)\}$ is computable. \square

Lemma 4. *In terms of this ad hoc computability structure, T is effectively determined.*

Proof. Since T is bounded, it suffices to show that T acts effectively on the generating set $\{f, e_1, e_2, \dots\}$ for H . We see at once that T acts effectively on the subspace H_0 generated by e_1, e_2, \dots . Thus all we need to show is that Tf is computable. Now

$$Tf = \sum_{n=1}^{\infty} 2^{-n} \alpha_n e_n.$$

Since $Tf \in H_0$, we can apply Lemma 1. The sequence of “Fourier coefficients” $\{2^{-n}\alpha_n\}$ is computable, and the sum $\sum 4^{-n}\alpha_n^2$ is effectively convergent (being dominated by $\sum 4^{-n}$). Hence Tf is computable. \square

We recall that the eigenvector of T corresponding to the eigenvalue 0 is Λ . Then we have:

Lemma 5. *The only constant multiple $c\Lambda$ of Λ which is ad hoc computable in H is the zero vector.*

Proof. We first note that if $c\Lambda$ is computable in H , then $|c|$ is a computable real. This follows from the Norm Axiom: for $\|c\Lambda\| = |c|$ must be computable if $c\Lambda$ is.

Now suppose $c\Lambda = \beta f + y$ is computable. Then the constant β and the vector y are computable. Furthermore, $f = \gamma\Lambda + \sum \alpha_n e_n$, where γ is the noncomputable real defined above. Since $y \in H_0$, the Λ -component of $\beta f + y$ is $\beta\gamma \cdot \Lambda$. Hence $c = \beta\gamma$, $|c| = |\beta|\gamma$ and $|c|$ is computable if and only if $\beta = 0$. \square

As noted above, the preliminary Eigenvector Theorem is an immediate consequence of Lemmas 3, 4, and 5.

6. The Eigenvector Theorem, Completed

This section treats two main topics. One, as stated in its heading, is the completion of the Eigenvector Theorem. The second—and related—topic is ad hoc computability, as introduced in Chapter 2, Section 7. We prove a number of general results about ad hoc computability. These results are of independent interest, since they can be used as tools in a variety of situations. One of these situations is the proof of the Eigenvector Theorem.

The ad hoc computability results (Lemmas 7 and 8) are given in their most general form. This entails a bit of extra work. In particular, it requires the inclusion of Lemma 6, whose slightly complicated proof is deferred until Section 7. We mention that, for the proof of the Eigenvector Theorem, Lemma 6 could be omitted. But the most general forms of Lemmas 7 and 8, which apply to an *arbitrary* computability structure on a Hilbert space, require Lemma 6.

Although most of the results in this section apply only to Hilbert space, the following holds for Banach spaces.

Lemma 6 (The Effective Independence Lemma). *Let X be an effectively separable Banach space with an effective generating set $\{e_n\}$. Then there exists an effective generating set $\{f_n\}$ for X whose elements are linearly independent over the real or complex numbers.*

As mentioned above, the proof of this lemma is deferred until Section 7.

Lemma 7. *Let H be a Hilbert space with an arbitrary computability structure imposed upon it. Let $\{e_n\}$ be an effective generating set for H . Then there exists a computable orthonormal basis $\{u_n\}$ for H .*

Proof. By Lemma 6, we can assume without loss of generality that the elements of $\{e_n\}$ are linearly independent. The rest is easy. Following the standard Gram-Schmidt process, we write:

$$v_0 = e_0,$$

$$v_n = e_n - \sum_{k=0}^{n-1} \frac{(e_n, v_k)}{(v_k, v_k)} v_k.$$

Then the vectors v_n are orthogonal, and by the Linear Forms Axiom the sequence $\{v_n\}$ is computable.

[Strictly speaking, the inductive definition of $\{v_n\}$ given above does not fit the format of the Linear Forms Axiom. Namely, this axiom requires that we have linear forms in the original computable sequence $\{e_n\}$. However, we can compute the v_n in terms of the e_n by means of back-substitution.]

Finally, since the sequence of norms $\{\|v_n\|\}$ is computable by the Norm Axiom (and since $v_n \neq 0$ for all n , by linear independence), we can write

$$u_n = v_n / \|v_n\|,$$

giving a computable orthonormal sequence $\{u_n\}$. \square

Remarks. Let us illustrate Lemma 7 as it applies to the Eigenvector Theorem. Begin with the effective generating set $\{f, e_1, e_2, \dots\}$, which gave the ad hoc computability structure defined in Section 5. We recall that $\{e_1, e_2, \dots\}$ is the standard computable orthonormal basis for the subspace H_0 , but that f is an “exotic” element which is not computable in the standard sense. Now the Gram-Schmidt process begins with f , giving $u_0 = f$ (since $\|f\| = 1$). But f is *not* orthogonal to e_1, e_2, \dots , and so the Gram-Schmidt process alters e_1, e_2, \dots , producing a new orthonormal sequence $\{u_0, u_1, u_2, \dots\}$. The sequence $\{u_n\}$ is not computable in the standard structure (for, as we have already seen, $u_0 = f$ is not computable in this sense). However, Lemma 7 implies that $\{u_n\}$ is computable in the ad hoc structure. Indeed, this is precisely the point of Lemma 7.

Now Lemma 8 below shows us how to translate results about the ad hoc computability structure into corresponding results for the standard structure. Although its proof is easy, it is a very useful result. In particular, it provides the final step in the proof of the Eigenvector Theorem.

Before turning to the technical details, we begin with a brief discussion of the content of Lemma 8. This Lemma asserts that, once we fix an orthonormal basis $\{u_k\}$ which is declared to be computable, the entire computability structure is fixed via conditions (i) and (ii) below. However, different computability structures can arise via the choice of different orthonormal bases. A basis which is computable in

terms of one computability structure might not be computable in terms of another. An example is the orthonormal basis $\{u_k\}$ above, which is ad hoc computable but not computable in the standard sense. Indeed, if the basis $\{u_k\}$ were computable in the standard sense, then by the Effective Density Lemma (Chapter 2, Section 5) the ad hoc computability structure would be the same as the standard one—which it clearly is not.

One final comment: Lemma 8 gives a representation theorem for any effectively separable computability structure on Hilbert space, and thus shows that all such structures are isomorphic. As we shall show in the next section, the corresponding statement for Banach spaces is false.

Lemma 8. *Consider any effectively separable computability structure—in the axiomatic sense—on a Hilbert space H . Then this structure has the following form. There is a computable orthonormal basis $\{u_k\}$ for H , and in terms of this basis: A sequence of vectors $\{x_n\}$ in H is computable if and only if*

i) *the double sequence of “Fourier coefficients” $\{c_{nk}\}$ of the x_n in terms of $\{u_k\}$ is computable, and*

ii) *the series $\sum_{k=0}^{\infty} |c_{nk}|^2$ converges effectively in k and n .*

Proof. Combine Lemma 1 from Section 5 with Lemma 7 above. \square

Proof of the Eigenvector Theorem, completed. We combine the preliminary theorem of Section 5 with Lemma 8. Let T be the operator on the ad hoc space H given in Section 5. Let u_0, u_1, u_2, \dots be the ad hoc computable orthonormal basis for H (so that the ad hoc computability structure is given via Lemma 8). Let e_0, e_1, e_2, \dots be the standard computable orthonormal basis for $L^2[0, 1]$.

Let \tilde{T} be the operator on $L^2[0, 1]$ which acts on the basis $\{e_0, e_1, e_2, \dots\}$ in the same way that T acts on $\{u_0, u_1, u_2, \dots\}$. That is, if

$$Tu = \sum_{k=0}^{\infty} b_{nk} u_k,$$

then

$$\tilde{T}e_n = \sum_{k=0}^{\infty} b_{nk} e_k.$$

What must we prove about \tilde{T} ? Here let us recall the essential properties of T , as proved in Section 5. We showed there that T is compact and self-adjoint, and that 0 is an eigenvalue of T of multiplicity one. We showed further that T is effectively determined in terms of the ad hoc computability structure on H . Finally we showed that, in terms of this ad hoc structure, none of the eigenvectors of T corresponding to $\lambda = 0$ is computable.

We claim that \tilde{T} , acting on the standard computability structure, shares all of the properties of T listed above. These properties are of two types:

- 1) geometric properties—compactness, self-adjointness, the values and multiplicities of the eigenvalues, and
- 2) computability properties.

For (1), since $\{u_k\}$ and $\{e_k\}$ are both orthonormal bases, the transition from $\{u_k\}$ to $\{e_k\}$ preserves all of the above mentioned geometric properties. For (2), Lemma 8 guarantees that the transition from $\{u_k\}$ to $\{e_k\}$ maps the ad hoc computability structure on H onto the natural computability structure on $L^2[0, 1]$.

Thus \tilde{T} , acting on the standard computability structure, is the operator promised in the main Eigenvector Theorem. The proof of that theorem is now complete. \square

One final question. What does the operator \tilde{T} actually look like? Well, \tilde{T} is just the operator T expressed in terms of the ad hoc computable orthonormal basis $\{u_k\}$. The basis $\{u_k\}$ is computed from the sequence $\{f, e_1, e_2, \dots\}$ via Lemma 7, and the formula for f , on which this computation depends, is given in Section 5. Beginning with these observations, an explicit presentation of \tilde{T} can be worked out. However, its form is rather a mess, and its intuitive meaning is completely hidden.

7. Some Results for Banach Spaces

For the first time in this chapter, we deal with computability structures on an arbitrary Banach space, rather than a Hilbert space. Our first result is the Effective Independence Lemma (Lemma 6 in Section 6), whose proof was postponed until this section. Our second result is a counterexample which shows that, on the Banach space l^1 , there exist effectively separable ad hoc computability structures which are not isometric to the standard structure. This contrasts with the situation for Hilbert space (Lemma 8 in Section 6).

Effective Independence Lemma. *Let X be an effectively separable Banach space with an effective generating set $\{e_n\}$. Then there exists an effective generating set $\{f_n\}$ for X whose elements are linearly independent over the real or complex numbers.*

Proof of lemma. We begin with the effective generating set $\{e_n\}$, which we already have. The sequence $\{f_n\}$ will consist of a subset of $\{e_n\}$, selected by an effective process to be linearly independent and have dense linear span in X . We emphasize that $\{f_n\}$ is *not* a subsequence of $\{e_n\}$, since the terms in $\{f_n\}$ may appear in a different order than they do in $\{e_n\}$.

Proof sketch. The idea behind the construction of $\{f_n\}$ is roughly as follows. The construction proceeds by induction. At the q -th stage, suppose we already have the first $(k + 1)$ elements f_0, \dots, f_k of the desired sequence $\{f_n\}$. Then we examine in turn all of the elements e_0, \dots, e_q which do not already belong to the set $\{f_0, \dots, f_k\}$, beginning with the e_i of smallest index and working upwards. For each e_i , we apply an effective “Test” (to be described below) which terminates in a finite number of

steps and leads to one of the two conclusions: (a) the set $\{f_0, f_1, \dots, f_k, e_i\}$ is linearly independent; (b) the vector e_i can be approximated to within a distance $< 2^{-q}$ by a linear combination of f_0, \dots, f_k . As soon as some e_i satisfies condition (a), we add e_i to $\{f_0, \dots, f_k\}$; that is, we set $f_{k+1} = e_i$. In case several e_i satisfy (a), we select only one: that with the smallest i . When no e_i satisfies (a), we do nothing. Then we go on to the $(q + 1)$ st stage.

By definition, the sequence $\{f_n\}$ so constructed is linearly independent. And by condition (b), any e_i which is forever omitted from $\{f_n\}$ can be arbitrarily closely approximated by linear combinations of the f_n ; hence e_i lies in the closed linear span of $\{f_n\}$. Thus, since by definition the linear span of $\{e_n\}$ is dense in X , so is the linear span of $\{f_n\}$. This completes the proof-sketch.

Details of the proof. We begin by giving a criterion for the linear independence of any string $\{z_1, \dots, z_k\}$ of computable vectors in X . This criterion will correspond to part (a) of the effective “Test” mentioned above.

By definition, the vectors $\{z_1, \dots, z_k\}$ are linearly dependent if and only if there exist (real/complex) scalars $\{\beta_1, \dots, \beta_k\}$, not all zero, such that $\beta_1 z_1 + \dots + \beta_k z_k = 0$. Multiplying through by a constant, we can assume that $\{\beta_1, \dots, \beta_k\}$ lies in the domain D between the unit sphere and the sphere of radius 2 in \mathbb{R}^k or \mathbb{C}^k . Now we give a simple recipe for approximating all of the points in D by (dyadic rationals/dyadic complex rationals).

For $m = 0, 1, 2, \dots$, let S_{mk} denote the set of all k -tuples $\{\beta_1, \dots, \beta_k\}$ of (rationals/complex rationals) whose denominators are 2^m and which satisfy

$$1 \leq |\beta_1|^2 + \dots + |\beta_k|^2 \leq 4.$$

Then, for each m and k , S_{mk} is a finite set, and there is an obvious procedure for listing S_{mk} for all m and k , effectively in m and k .

Independence Criterion. *The vectors $\{z_1, \dots, z_k\}$ are linearly independent if and only if the following condition holds. For some $m \geq 2k$:*

$$\min \{ \|\beta_1 z_1 + \dots + \beta_k z_k\| : \{\beta_1, \dots, \beta_k\} \in S_{mk} \} > 2^{-m} \cdot (\|z_1\| + \dots + \|z_k\|).$$

Note. We observe that this criterion involves the evaluation of $\|\beta_1 z_1 + \dots + \beta_k z_k\|$ at only a finite number of points, namely the points in the set S_{mk} . By the Linear Forms and Norm Axioms, $\|\beta_1 z_1 + \dots + \beta_k z_k\|$ is computable, effectively in k , the β_i , and z_i . Similarly for $(\|z_1\| + \dots + \|z_k\|)$.

Proof of criterion. For the “if” part. Suppose that the criterion holds, but that $\{z_1, \dots, z_k\}$ are linearly dependent. We must derive a contradiction. Since $\{z_1, \dots, z_k\}$ are dependent, there exists a point $\{\gamma_1, \dots, \gamma_k\}$ on the sphere of radius $3/2$ in \mathbb{R}^k or \mathbb{C}^k (i.e. with $|\gamma_1|^2 + \dots + |\gamma_k|^2 = 9/4$) such that $\gamma_1 z_1 + \dots + \gamma_k z_k = 0$. Now for each γ_i , there exists a dyadic value β_i , with denominator 2^m , such that

$$|\gamma_i - \beta_i| \leq 2^{-m}.$$

[Actually, in the real case we can get $|\gamma_i - \beta_i| \leq (1/2)2^{-m}$; and in the complex case we can get a distance of $(\sqrt{2}/2)2^{-m}$.] Thus, since $\gamma_1 z_1 + \dots + \gamma_k z_k = 0$,

$$\|\beta_1 z_1 + \dots + \beta_k z_k\| \leq \sum_{i=1}^k |\gamma_i - \beta_i| \cdot \|z_i\| \leq 2^{-m}(\|z_1\| + \dots + \|z_k\|).$$

This contradicts the inequality given in the criterion.

For the “only if” parts. If $\{z_1, \dots, z_k\}$ are linearly independent, then over the entire domain $D = \{1 \leq |\beta_1|^2 + \dots + |\beta_k|^2 \leq 4\}$ (without any reference to dyadic points), the minimum value δ ,

$$\delta = \min \{\|\beta_1 z_1 + \dots + \beta_k z_k\| : \{\beta_1, \dots, \beta_k\} \in D\}$$

satisfies $\delta > 0$.

Now simply take any m large enough so that $2^{-m}(\|z_1\| + \dots + \|z_k\|) < \delta$, and we see that the criterion does hold, as desired. This proves the validity of the criterion. \square

Construction of $\{f_n\}$. Now we begin the construction of the sequence $\{f_n\}$. At this point we must take pains to work within the axioms for computability on a Banach space.

We start with the given effective generating set $\{e_n\}$. Then we sweep out the set of all finite (rational/complex rational) linear combinations of the e_n ; this can be done in an effective way by using one of the standard recursive enumerations of all finite sequences of integers. By the Linear Forms Axiom, this yields a computable sequence $\{p_n\}$ in X which consists of all finite (rational/complex rational) linear combinations of the e_n . Then by the Norm Axiom, the sequence of norms $\{\|p_n\|\}$ is a computable sequence of reals.

We return now to the inductive definition of the sequence $\{f_n\}$, as described briefly in the proof-sketch above. We recall that at stage q we had selected $(k + 1)$ linearly independent vectors $\{f_0, \dots, f_k\}$ for $\{f_n\}$, and that we would then apply a “Test” to all of the elements e_0, \dots, e_q not already in the set $\{f_0, \dots, f_k\}$. Here is the Test:

The Test. The test has two parts, (a) and (b), between which we alternate, switching back and forth until a termination is reached. The test is applied to an element e_i from the set e_0, \dots, e_q .

Both of the parts (a) and (b) below involve strict inequalities (“ $>$ ” and “ $<$ ” respectively). Hence, IF either of these inequalities holds, a finite amount of calculation (involving a sufficiently good rational approximation) will suffice to confirm it. The fact that this procedure halts will be proved below. As we shall see, (a) and (b) are not mutually exclusive. Indeed, this is why an effective decision procedure is possible.

Part (a). With reference to the “Independence Criterion” above: For $m = 2k, 2k + 1, 2k + 2, \dots$, compute all of the values $\|\beta_0 f_0 + \dots + \beta_k f_k + \beta_{k+1} e_i\|$ for the finite set of points $\{\beta_0, \dots, \beta_{k+1}\} \in S_{m, k+2}$. Using rational approximations as discussed

above, test whether:

$$\begin{aligned} \min \{ \|\beta_0 f_0 + \cdots + \beta_k f_k + \beta_{k+1} e_i\| : \{\beta_0, \dots, \beta_{k+1}\} \in S_{m, k+2} \} \\ > 2^{-m} (\|f_0\| + \cdots + \|f_k\| + \|e_i\|). \end{aligned}$$

If this ever happens, for any m , cry “Halt!” and declare that “(a) holds: $\{f_0, \dots, f_k, e_i\}$ are linearly independent.”

Part (b). (This part is disappointingly unsubtle.) Scan through all of the (rational/complex rational) linear combinations of f_0, \dots, f_k until such a combination $\beta_0 f_0 + \cdots + \beta_k f_k$ is found for which

$$\|e_i - (\beta_0 f_0 + \cdots + \beta_k f_k)\| < 2^{-q}.$$

If this happens, cry “Halt!” and declare “(b) holds: e_i can be approximated to within a distance of 2^{-q} by a linear combination of f_0, \dots, f_k .”

Proof that “The Test” halts. By hypothesis, the vectors $\{f_0, \dots, f_k\}$ are linearly independent. Therefore, either (A) the vectors $\{f_0, \dots, f_k, e_i\}$ are independent, or (B') e_i is a linear combination of f_0, \dots, f_k . By the “Independence Criterion”, (A) is equivalent to (a) above. Now (B') is stronger than (b), but it suffices that (B') implies (b). Thus either (a) or (b) (or both) holds.

We still have to verify that one of the processes (a) or (b) halts. But this is easy. Both processes involve strict inequalities, “ $>$ ” and “ $<$ ” respectively. Thus IF either of the inequalities (a) or (b) is true, a finite amount of calculation will verify it, and the process will eventually halt. Since at least one of (a) and (b) is true, the process does halt, as desired. \square

Now the rest of the proof is nearly identical to that given in the proof-sketch above. At stage q , we apply “The Test” to e_0, \dots, e_q (omitting those e_i already in $\{f_0, \dots, f_k\}$). When the first e_i satisfies part (a) of the test, we add e_i to $\{f_0, \dots, f_k\}$, setting $f_{k+1} = e_i$. Then we stop stage q . If no e_i satisfies part (a), we do nothing in stage q . Then we go on to stage $(q + 1)$.

As explained in the proof-sketch above, the sequence $\{f_n\}$ is linearly independent and its closed linear span contains $\{e_n\}$; hence $\{f_n\}$ is dense in X .

Finally we must show that $\{f_n\}$ is computable in X . Now the construction of $\{f_n\}$ involved a recursive process, depending on the computable sequence of real numbers $\{\|p_n\|\}$. Thus f_n has the form $f_n = e_{a(n)}$ for a recursive function $a: \mathbb{N} \rightarrow \mathbb{N}$. By the Composition Property, proved as a consequence of the Axioms in Chapter 2, $\{f_n\}$ is computable. \square

A counterexample for Banach spaces

Now that we have illustrated the usefulness of ad hoc computability structures, it seems natural to pose the following question. Are all such structures related via (not necessarily computable) isometries? We recall the definition:

Definition. Let X be a Banach space. A linear transformation $U: X \rightarrow X$ is called an *isometry* if U is onto and distance preserving (i.e. U is onto and $\|Ux\| = \|x\|$).

Of course, if U is an isometry, then U^{-1} exists and is an isometry.

At this point, we dispose of a triviality. Certain non-isometric mappings can give the same computability structure as an isometric mapping. For example, if U is an isometry and $k > 0$ is a computable real, then kU gives the same structure as U . Obviously such examples add nothing new to our problem, which as precisely formulated is:

Question. Are all effectively separable ad hoc computability structures \mathcal{S}_1 related to the standard structure \mathcal{S}_0 in the following way? There exists a (not necessarily computable) isometry U such that

$$\{f_n\} \in \mathcal{S}_1 \quad \text{if and only if } f_n = Ug_n \text{ for some } \{g_n\} \in \mathcal{S}_0.$$

All of the ad hoc computability structures which we have used so far have been based on non-computable isometries. Thus the two examples given in Chapter 2, Section 7 involved (1) multiplication by a non-computable complex constant c with $|c| = 1$, and (2) the translation $f(x) \rightarrow f(x + \alpha)$, where α is a noncomputable real. Both of these transformations are isometries. The more recondite example in Sections 5 and 6 involved an ad hoc computability structure on a Hilbert space. We proved in Lemma 8 above that any ad hoc computability structure on an effectively separable Hilbert space is isometric to the natural structure.

Suppose, then, that we turn our attention to Banach spaces. Here the situation is different.

Example. There exist ad hoc structures on l^1 which are not isometric to the standard structure.

To give such an example is not trivial. [The difficulty, of course, lies in satisfying the Norm Axiom. With the isometries—which preserved norms—this was no problem.] Since this example is used nowhere else in the book, we shall present it in a rather terse fashion. We remark that some of the steps are similar to those used in the Preliminary Eigenvector Theorem of Section 5.

Proof. We consider the real Banach space l^1 , and let e_0, e_1, \dots denote the unit vectors, $e_n = \{0, 0, \dots, 0, 1, 0, \dots\}$ with a “1” in the n -th place. For typographical clarity we denote e_0 by Λ .

As usual, let $a: \mathbb{N} \rightarrow \mathbb{N}$ be a one to one recursive function generating a recursively enumerable nonrecursive set A . We assume that $0 \notin A$. Let

$$\alpha_n = 2^{-a(n)},$$

and set

$$\gamma = 1 - \sum_{n=1}^{\infty} \alpha_n.$$

Then γ is a noncomputable real. Let

$$f = \gamma\Lambda + \sum_{n=1}^{\infty} \alpha_n e_n.$$

Then, by definition of γ , $\|f\| = 1$.

Let l_0^1 denote the subspace of l^1 spanned by $\{e_1, e_2, \dots\}$.

We define the ad hoc computability structure on l^1 by specifying that

$$\{f, e_1, e_2, e_3, \dots\}$$

is an effective generating set.

Equivalently, a sequence $\{x_n\}$ of vectors in l^1 is *ad hoc computable* if:

$$x_n = \beta_n f + y_n$$

where $\{\beta_n\}$ is a computable sequence of real numbers, and $\{y_n\}$ is a computable sequence of vectors (in the standard sense) in l_0^1 .

Lemma. *This is a computability structure.*

Proof. The Linear Forms and Limit Axioms are clear. Only the Norm Axiom requires proof. That is, we must show that $\{\|x_n\|\}$ is computable if $\{x_n\}$ is. Fix n , and consider a single computable vector x . It will be clear that the procedure which follows is effective in n .

We have:

$$x = \beta f + y,$$

where β is a computable real number and y is computable in l_0^1 .

Let $y = \{\theta_1, \theta_2, \dots\}$, $\theta_i \in \mathbb{R}$.

Since y is computable in l^1 , there exists a recursive function $e(N)$ such that

$$\sum_{k=e(N)}^{\infty} |\theta_k| \leq 2^{-N} \quad \text{for all } N.$$

To compute $\|x\|$ to within 2^{-N} , we use the following recipe:

Compute $\alpha_0, \alpha_1, \dots, \alpha_{e(N)}$

Compute $\beta\alpha_0 + \theta_0, \dots, \beta\alpha_{e(N)} + \theta_{e(N)}$.

Now

$$\|x\| = \sum_{k=1}^{\infty} |\beta\alpha_k + \theta_k| + |\beta|\gamma,$$

and to within an error of 2^{-N} (gotten by dropping the “tail” of $\{\theta_k\}$):

$$\|x\| = \sum_{k=1}^{e(N)} |\beta\alpha_k + \theta_k| + \sum_{k=e(N)+1}^{\infty} |\beta|\alpha_k + |\beta|\gamma.$$

The last two terms in the above displayed formula are not computable. However we claim that their sum

$$|\beta| \left(\gamma + \sum_{k=e(N)+1}^{\infty} \alpha_k \right)$$

is computable.

To show this, we reason as follows. Firstly, $|\beta|$ is computable. We know that

$$\gamma + \sum_{k=1}^{\infty} \alpha_k = 1,$$

so

$$\gamma + \sum_{k=e(N)+1}^{\infty} \alpha_k = 1 - \sum_{k=1}^{e(N)} \alpha_k. \quad \square$$

Lemma. *This structure is not isometric to the standard one.*

Proof. We recall the definition of an “extremal point” on the closed unit ball B of a Banach space. (Recall that $B = \{x: \|x\| \leq 1\}$.) We say that a vector $u \in B$ is *extremal* if there do not exist distinct vectors $v, w \in B$ and a constant $c, 0 < c < 1$, such that $u = cv + (1 - c)w$. Clearly any isometry preserves these extremal points.

It is well-known and easy to show that the only extremal points in l^1 are $\pm e_k$.

Now the standard computability structure on l^1 has an effective generating set consisting entirely of extremal points (namely the e_k , including $\Lambda = e_0$).

Thus, since isometries preserve extremal points, if the ad hoc structure were isometric to the standard one, the ad hoc structure would also have an effective generating set consisting entirely to extremal points. We now show that this cannot happen.

Suppose otherwise. The ad hoc effective generating set must contain some element z with a nonzero Λ -component. As we have seen, if z is extremal, then $z = \pm \Lambda$. On the other hand z —since it comes from an effective generating set—must be ad hoc computable. This is impossible in view of the following:

Lemma. *The only multiple of Λ which is ad hoc computable is $0 \cdot \Lambda$.*

Proof. Virtually identical to the proof of Lemma 5 in Section 5. \square

This completes the proof of the previous lemma, and hence also of the main result. \square