

Operator Algebras

Nigel Higson and John Roe

Department of Mathematics, The Pennsylvania State University, University Park, PA 16802, USA

1 The beginnings of operator theory

We can ask two basic questions about any equation, or system of equations: is there a solution, and, if there is, is it unique? Experience with finite systems of linear equations indicates that the two questions are interconnected. Consider for instance the equations

$$\begin{aligned}2x + 3y - 5z &= a \\ x - 2y + z &= b \\ 3x + y - 4z &= c\end{aligned}$$

Notice that the left-hand side of the third equation is the sum of the left-hand sides of the first two. As a result, no solution to the system exists unless $a + b = c$. But if $a + b = c$, then any solution of the first two equations is also a solution of the third; and in any linear system involving more unknowns than equations, solutions, when they exist, are never unique. In the present case, if (x, y, z) is a solution, then so is $(x + t, y + t, z + t)$, for any t . Thus the same phenomenon (a linear relation among the equations) which prevents the system from admitting solutions in some cases, also prevents solutions from being unique in other cases.

To make the relation between existence and uniqueness of solutions more precise, consider a general system of linear equations of the form

$$\begin{aligned}k_{11}u_1 + k_{12}u_2 + \cdots + k_{1n}u_n &= f_1 \\ k_{21}u_1 + k_{22}u_2 + \cdots + k_{2n}u_n &= f_2 \\ &\vdots \\ k_{n1}u_1 + k_{n2}u_2 + \cdots + k_{nn}u_n &= f_n\end{aligned}$$

consisting of n equations in n unknowns. The scalars k_{ji} form a matrix of coefficients and the problem is to solve for the u_i in terms of the f_j . The general theorem illustrated by our particular numerical example above is that the number of linear conditions among the f_j required for a solution to exist is equal to the number of arbitrary

constants appearing in the general solution when a solution does exist. To use a more elaborate vocabulary, the dimension of the kernel of the matrix $K = \{k_{ji}\}$ is equal to the dimension of its cokernel. In the example, these numbers are both 1.

A little more than 100 years ago Ivar Fredholm made a study of *integral equations* of the type

$$u(y) - \int k(y, x)u(x) dx = f(y).$$

These arose from questions in theoretical physics, and the problem is to solve for the function u in terms of the function f . Since an integral can be thought of as a limit of finite sums, Fredholm's equation is an infinite-dimensional counterpart of the finite-dimensional linear systems considered above, in which vectors with n components are replaced by functions with values at infinitely many different points x . (Strictly speaking, Fredholm's equation is analogous to a matrix equation of the type $u - Ku = f$ rather than $Ku = f$. The altered form of the left-hand side has no effect on the overall behavior of the matrix equation, but it does considerably alter the behavior of the integral equation. As we shall see, Fredholm was fortunate to work with a class of equations whose behavior mirrors very closely that of matrix equations.)

A very simple example is

$$u(y) - \int_0^1 u(x) dx = f(y).$$

To solve this equation, it helps to observe that the quantity $\int_0^1 u(x) dx$, when thought of as a function of y , is a *constant*. Thus in the homogeneous case ($f \equiv 0$), the only possible solutions for $u(y)$ are the constant functions. On the other hand, for a general function f , solutions exist if and only if the single linear condition $\int_0^1 f(y) dy = 0$ is satisfied. So in this example the dimension of the kernel and the dimension of the cokernel are both 1. Fredholm set out on a systematic exploration of the analogy between matrix theory and integral equations which this example suggests. He was able to prove that, for equations of his type, the dimensions of the kernel and of the cokernel are always finite and equal.

Fredholm's work sparked the imagination of David Hilbert, who made a detailed study of the *integral operators* which transform $f(y)$ into

$\int k(y, x)u(x) dx$ in the special case where the real-valued function k is *symmetric*, meaning $k(x, y) = k(y, x)$. The finite-dimensional counterpart is the theory of real symmetric matrices. Now if K is such a matrix, then a standard result from linear algebra asserts that there is an orthonormal basis consisting of EIGENVECTORS for K , or equivalently that there is a *unitary* matrix U such that $U^{-1}TU$ is diagonal. (*Unitary* means that U is invertible and preserves the lengths of vectors: $\|Uv\| = \|v\|$ for all vectors v .) Hilbert obtained an analogous theory for all symmetric integral operators. He showed that there exist functions $u_1(y), u_2(y), \dots$ and real numbers $\lambda_1, \lambda_2, \dots$ such that

$$\int k(y, x)u_n(x) dx = \lambda_n u_n(y).$$

Thus $u_n(y)$ is an *eigenfunction* for the integral operator, with eigenvalue λ_n .

In most cases it is hard to calculate u_n and λ_n explicitly, but calculation *is* possible when $k(x, y) = \phi(x - y)$ for some periodic function ϕ . If the range of integration in $[0, 1]$ and the period of ϕ is 1, then the eigenfunctions are $\cos(2k\pi y)$, $k = 0, 1, 2, \dots$, and $\sin(2k\pi y)$, $k = 1, 2, \dots$. In this case the theory of FOURIER SERIES tells us that a general function $f(y)$ on $[0, 1]$ can be expanded as the sum of a series $\sum (a_k \cos 2k\pi y + b_k \sin 2k\pi y)$ of cosines and sines. Hilbert showed that, in general, there is an analogous expansion

$$f(y) = \sum a_n u_n(y)$$

in terms of the eigenfunctions for *any* symmetric integral operator. In other words, the eigenfunctions form a *basis*, just as in the finite-dimensional case. Hilbert's result is now called the *spectral theorem* for symmetric integral operators.

1.1 From integral equations to functional analysis

Hilbert's theorem led to an explosion of activity, since integral operators arise in many different areas of mathematics (including for example the Dirichlet problem in partial differential equations and the representation theory of compact groups). It was soon recognized that these operators are best viewed as linear transformations

on the HILBERT SPACE of all functions $u(y)$ such that $\int |u(y)|^2 dy < \infty$. Such functions are called *square-integrable*, and the collection of all of them is denoted $L^2[0, 1]$.

With the important concept of Hilbert space available, it became convenient to examine a much broader range of operators than the integral operators initially considered by Fredholm and Hilbert. Since Hilbert space is both a vector space and a metric space, it was natural to look first at linear operators from a Hilbert space to itself which are continuous—these are usually called *bounded* linear operators. The analogue of the symmetry condition $k(x, y) = k(y, x)$ on integral operators is the condition that a bounded linear operator T be *self-adjoint*, which is to say that $\langle Tu, v \rangle = \langle u, Tv \rangle$ for all vectors u and v in the Hilbert space (the angle brackets denote the inner product). A simple example of a self-adjoint operator is the *multiplication operator* by a real-valued function $m(y)$; this is the operator M defined by the formula $(Mu)(y) = m(y)u(y)$. (The finite-dimensional counterpart to a multiplication operator is a diagonal matrix K , which multiplies the j th component of the vector by the matrix entry k_{jj} .)

The *spectrum* of an operator T is the set of complex numbers λ for which the operator $T - \lambda I$ does not have a bounded inverse (here I is the identity operator on Hilbert space). In finite dimensions the spectrum is precisely the set of eigenvalues, but in infinite dimensions this is not always so. In fact, whereas every symmetric matrix has at least one eigenvalue, a self-adjoint operator need have none at all. As a result of this, the spectral theorem for bounded self-adjoint operators is phrased not in terms of eigenvalues but in terms of the spectrum. One way of formulating it is to state that any self-adjoint operator T is unitarily equivalent¹ to a multiplication operator $(Mu)(y) = m(y)u(y)$, where the closure of the range of the function $m(y)$ is the spectrum of T . This generalizes the statement that any real symmetric matrix is unitarily equivalent to a diagonal matrix with the eigenvalues along the diagonal.

¹Just as in the finite-dimensional case, a *unitary* is an invertible operator that preserves the lengths of vectors, and *unitary equivalence* means that $M = U^{-1}TU$.

1.2 The mean ergodic theorem

A beautiful application of the spectral theorem was found by John von Neumann. Imagine a checkerboard on which are distributed a certain number of checkers. Imagine that for each square there is designated a “successor” square (in such a way that no two squares have the same successor), and that every minute the checkers are rearranged by moving each one to its successor square. Now focus attention on a single square and record every minute whether or not there is a piece on the square, so producing a succession of readings R_1, R_2, R_3, \dots like this:

00100110010110100100...

We might expect that over time, the average number of positive readings $R_j = 1$ will converge to the number of pieces on the board divided by the number of squares. If the rearrangement rule is not complicated enough, then this will not happen. For example, in the most extreme case, if the rule designates each square as its own successor, then the readout will be either 00000... or 11111..., depending on whether or not we chose a square with a piece on it to begin with. But if the rule is sufficiently complicated then the “time average” $\frac{1}{n} \sum_{j=1}^n R_j$ will indeed converge to the number of pieces on the board divided by the number of squares, as expected.

The checkerboard example is elementary, since in fact the only “sufficiently complicated” rules in this finite case are CYCLIC PERMUTATIONS of the squares of the board, and thus *all* the squares move past our observation post in succession. However, there are related examples where only a small part of the data is actually observed. For instance, replace the set of squares on a checkerboard with the set of points on a circle, and in place of the checkers, imagine that a subset S of a circle is marked as occupied. Let the rearrangement rule be the rotation of points on the circle through some irrational number of degrees. Stationed at a point x of the circle, we record whether x belongs to S , the first rotated copy of S , the second rotated copy of S , and so on to obtain a sequence of 0 or 1 readings as before. One can show that (for nearly every x) the time average of our observations will converge to the proportion of the circle occupied by S .

Similar questions about the relationship between

time and space averages had arisen in thermodynamics and elsewhere, and the expectation that time and space averages should agree when the rearrangement rule is sufficiently complex became known as the *ergodic hypothesis*.

Von Neumann brought operator theory to bear on this question in the following way. Let H be the Hilbert space of functions on the squares of the checkerboard, or the Hilbert space of square-integrable functions on the circle. The rearrangement rule gives rise to a unitary operator U on H by means of the formula

$$(Uf)(y) = f(\phi^{-1}(y)),$$

where ϕ is the function describing the rearrangement. Von Neumann’s ergodic theorem asserts that if no nonconstant function in H is fixed by U (this is one way of saying that the rearrangement rule is “sufficiently complicated”), then for every function $f \in H$, the limit $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{j=1}^n U^j f$ exists and is equal to the constant function whose value everywhere is the average value of f . (To apply this to our examples, take $f(x)$ to be the function which is 1 if the point x is occupied and 0 otherwise.)

Von Neumann’s theorem can be deduced from a spectral theorem for unitary operators which is analogous to the spectral theorem for self-adjoint operators. Every unitary operator can be reduced to a multiplication operator, not by real-valued functions but by functions whose values are complex numbers of absolute value 1. The key to the proof then becomes a statement about complex numbers of absolute value 1: if z is such a complex number, different from 1, then the expression $\frac{1}{n} \sum_{j=1}^n z^j$ approaches zero as $n \rightarrow \infty$. This in turn is easily proved using the formula for the sum of a geometric series, $\sum_{j=1}^n z^j = z(1 - z^n)/(1 - z)$.

1.3 Operators and quantum theory

Von Neumann realized that Hilbert spaces and their operators provide the correct mathematical tools to formalize the laws of quantum mechanics, introduced in the 1920’s by Heisenberg and Schrödinger.

The *state* of a physical system at an instant of time is the list of all the information needed to determine its future behavior. If, for instance, the

system consists of a finite number of particles, then classically its state consists of the list of the position and momentum vectors of all the constituent particles. In contrast, in von Neumann's formulation of quantum mechanics, to each physical system there is associated a Hilbert space H , and a state of the system is represented by a unit vector u in H . Vectors which differ only by multiplication with a scalar determine the same state.

Associated to each observable quantity (perhaps the total energy of the system, or the momentum of one particle within the system) is a self-adjoint operator Q on H whose spectrum is the set of all observed values of that quantity (hence the origin of the term "spectrum"). States and observables are related as follows: when a system is in the state described by a unit vector $u \in H$, the *expected value* of the observable quantity corresponding to a given self-adjoint operator Q is the inner product $\langle Qu, u \rangle$. This may not be a value which is ever actually measured, but it is rather the average of values which are obtained from many repeated experiments with system in the given state u . The relation between states and observables reflects the paradoxical behavior of quantum mechanics: it is possible, and in fact typical, for a system to exist in a "superposed" state, under which repeated identical experiments produce distinct outcomes. A measurement of an observable quantity will produce a determinate outcome if and only if the system is in a state which is an eigenvector for the operator associated to that quantity.

A distinctive feature of quantum theory is that the operators associated to different observables typically do not commute with one another. If two operators do not commute then they will typically have no eigenvectors in common, and, as a result, simultaneous measurements of two different observables will typically not result in determinate values for both of them. A famous example is provided by the operators P and Q associated to the position and momentum of a particle moving along a line. They satisfy the *Heisenberg commutation relation*

$$QP - PQ = i\hbar I,$$

where \hbar is a certain physical constant. (This is an instance of a general principle which relates the non-commutativity of observables in quantum me-

chanics to the *Poisson bracket* of the corresponding observables in classical mechanics.) As a result, it is impossible for the particle simultaneously to have a determinate momentum and position. This is the *uncertainty principle*.

It turns out that there is an essentially unique way of representing the Heisenberg commutation relation using self-adjoint operators on Hilbert space: the Hilbert space H must be $L^2(\mathbb{R})$; the operator P must be $-i\hbar d/dx$; and the operator Q must be multiplication by x . This theorem allows one to determine explicitly the observable operators for simple physical systems. For example, in a system consisting of a particle on a line subject to a force directed toward the origin which is proportional to the distance from the origin (as if the particle were attached to a spring, anchored at the origin), the operator for total energy is

$$E = -\frac{\hbar^2}{2m} \frac{d^2}{dx^2} + \frac{k}{2} x^2,$$

where k is a constant which determines the overall strength of the force. The spectrum of this operator is the set

$$\{(n + \frac{1}{2})\hbar(k/m)^{1/2} : n = 0, 1, 2, \dots\}.$$

These are therefore the possible values for the total energy of the system. Notice that the energy can assume only a discrete set of values—another characteristic and fundamental feature of quantum theory.

Another important example is the operator of total energy for the hydrogen atom. Like the operator above, this may be realized as a certain explicit partial differential operator. It may be shown that its eigenvalues form a sequence proportional to $\{-1, -1/4, -1/9, \dots\}$. A hydrogen atom, when disturbed, may release a photon, resulting in a drop in its total energy. The released photon will have energy equal to the difference in energies of the initial and final states of the atom, and hence proportional to $1/n^2 - 1/m^2$. When light from hydrogen is passed through a prism or diffraction grating, bright lines are indeed observed at wavelengths corresponding to these possible energies. Spectral observations of this sort provide experimental confirmation for quantum mechanical predictions.

So far we have discussed states of a quantum system only at a single instant of time. Quantum

systems however evolve in time just as classical systems do: to describe this evolution we need a law of motion. The time evolution of a quantum system is represented by a family of unitary operators $U_t: H \rightarrow H$, parameterized by the real numbers. If the system is in an initial state u , it will be in the state $U_t u$ after t units of time. Because the passage of s units of time followed by t further units is the same as the passage of $s+t$ units, the unitary operators U_t satisfy the *group law* $U_s U_t = U_{s+t}$. An important theorem of Marshall Stone asserts that there is a one-to-one correspondence between unitary groups $\{U_t\}$ and self-adjoint operators E given by the formula

$$iE = \left(\frac{dU_t}{dt} \right)_{t=0} = \lim_{t \rightarrow 0} \frac{1}{t} (U_t - I).$$

The quantum law of motion is that the *generator* E corresponding in this way to time evolution is the operator associated to the observable “total energy”. When E is realized as a differential operator on a Hilbert space of functions (as in the examples above), this statement becomes a differential equation, the *Schrödinger equation*.

1.4 The GNS construction

The time evolution operators U_t of quantum mechanics satisfy the law $U_s U_t = U_{s+t}$. More generally we define a *unitary representation* of a GROUP G to be a family of unitary operators U_g , one for each $g \in G$, satisfying the law $U_{g_1 g_2} = U_{g_1} U_{g_2}$ for all $g_1, g_2 \in G$. Originally introduced by Georg Frobenius as a tool for the study of finite groups, representation theory has become indispensable in mathematics and physics wherever the symmetries of a system must be taken into account.

If U is a unitary representation of G and v is a vector, then $\sigma: g \mapsto \langle U_g v, v \rangle$ is a function on G . The law $U_{g_1 g_2} = U_{g_1} U_{g_2}$ implies that σ has an important positivity property, namely

$$\sum_{g_1, g_2 \in G} \overline{a_{g_1}} a_{g_2} \sigma(g_1^{-1} g_2) = \left\| \sum a_g U_g v \right\|^2 \geq 0,$$

for any scalars $a_g \in \mathbb{C}$. A function on G with this positivity property is said to be *positive-definite*. Conversely, from a positive-definite function one can build a unitary representation. This *GNS construction* (in honor of Israel Gelfand, Mark

Naimark and Irving Segal) begins by considering the group elements themselves as basis vectors in an abstract vector space. We can attempt to define an inner product on this vector space by means of the formula

$$\langle g_1, g_2 \rangle = \sigma(g_1^{-1} g_2).$$

The resulting object may differ from a genuine Hilbert space in two respects. First, there may be nonzero vectors whose length, as measured by the inner product, is zero (although the hypothesis that σ is positive-definite does rule out the possibility that there might be vectors of *negative* length). Second, the COMPLETENESS AXIOM of Hilbert space theory may not be satisfied. However, there is a “completion” procedure which fixes both these deficiencies. Applied in the present case, it produces a Hilbert space H_σ which carries a unitary representation of G .

Versions of the GNS construction arise in several areas of mathematics. They have the advantage that the functions on which the constructions are based are easy to manipulate. For instance, convex combinations of positive-definite functions are again positive-definite, and this allows geometrical methods of to be applied to the study of representations.

1.5 Determinants and traces

The original works of Fredholm and Hilbert borrowed heavily from traditional concepts of linear algebra, and in particular the theory of DETERMINANTS. In view of the complicated definition of the determinant even for finite matrices, it is perhaps not surprising that the infinite-dimensional situation presented extraordinary challenges. Very soon, much simpler alternative approaches were found which avoided determinants altogether. But it is interesting to note that the determinant, or to be more exact the related notion of the trace, has played an important role in recent developments on which we will report later in this article.

The *trace* of an $n \times n$ matrix is the sum of its diagonal entries. Like the determinant, the trace has the property that the trace of BAB^{-1} is equal to the trace of A , for any invertible matrix B . In fact, the trace is related to the determinant by the formula $\det(\exp(A)) = \exp(\text{tr}(A))$ (because of the

invariance properties of trace and determinant it is enough to check this for diagonal matrices, where it is easy). In infinite dimensions the trace need not make sense since the sum of the diagonal entries of an $\infty \times \infty$ matrix may not converge. (The trace of the identity operator is a case in point: the diagonal entries are all 1, and if there are infinitely many of them, then their sum is not well-defined.) One way to address this problem is to limit oneself to operators for which the sum *is* well-defined. An operator T is said to be of *trace class* if for every two sequences $\{u_j\}$ and $\{v_j\}$ of pairwise orthogonal vectors of length 1, the sum $\sum_{j=1}^{\infty} \langle Tu_j, v_j \rangle$ is absolutely convergent. A trace-class operator T has a well-defined and finite trace, namely the sum $\sum_{j=1}^{\infty} \langle Tu_j, u_j \rangle$ (which is independent of the choice of orthonormal basis $\{u_j\}$).

Integral operators such as those appearing in Fredholm's equation provide natural examples of trace class operators. If $k(y, x)$ is a smooth function then the operator $Tu(y) = \int k(y, x)u(x)dx$ is of trace class, and its trace is equal to $\int k(x, x)dx$ —the “sum” of the diagonal elements of the “continuous matrix” k .

2 Von Neumann algebras

The *commutant* of a set S of bounded linear operators on a Hilbert space H is the collection S' of all operators on H which commute with every operator in the set S . The commutant of any set is an *algebra* of operators on H . That is, if T_1 and T_2 are in the commutant, then so are T_1T_2 and any linear combination $a_1T_1 + a_2T_2$.

Recall that a *unitary representation* of a group G on a Hilbert space H is a collection of unitary operators U_g , labeled by elements of G , which compose according to the law $U_{g_1g_2} = U_{g_1}U_{g_2}$. A *von Neumann algebra* is any algebra of operators on a complex Hilbert space H which is the commutant of some unitary representation of a group on H . Every von Neumann algebra is closed under adjoints and under limits of nearly every sort. For example, it is closed under pointwise limits: if $\{T_n\}$ is a sequence of operators in a von Neumann algebra M , and if $T_nv \rightarrow Tv$, for every vector $v \in H$, then $T \in M$.

It is easy to check that every von Neumann algebra M is equal to its own double commutant M''

(the commutant of the commutant of M). Von Neumann proved that if a self-adjoint algebra M of operators is closed under pointwise limits then M is equal to the commutant of the group of unitary operators in its commutant, and is therefore a von Neumann algebra.

2.1 Decomposing representations

Let $g \rightarrow U_g$ be a unitary representation of a group G on a Hilbert space H . If a closed subspace H_0 of H is mapped into itself by all the operators U_g , then it is said to be an *invariant subspace* for the representation. If H_0 is invariant, then since the operators U_g map H_0 to itself, their restrictions to H_0 constitute another representation of G , called a *subrepresentation* of the original.

A subspace H_0 is invariant for a representation, and so determines a subrepresentation, if and only if the orthogonal projection operator $P: H \rightarrow H_0$ belongs to the commutant of that representation. This points to a close connection between representation theory and von Neumann algebras. In fact von Neumann algebra theory can be thought of as a study of the ways in which unitary representations can be decomposed into subrepresentations.

A representation is *irreducible* if it has no non-trivial invariant subspace. A representation that does have a non-trivial invariant subspace H_0 can be divided into two subrepresentations—those associated to H_0 and to its orthogonal complement H_0^\perp . Unless both the representations H_0 and H_0^\perp are irreducible, we will be able to divide one or both of them into still smaller pieces by repeating the process that was just carried out for H . If the initial Hilbert space H is finite-dimensional, then continuing in this way we will eventually decompose it into irreducible subrepresentations. In the language of matrices, we will obtain a basis for H with respect to which all the operators in the group are simultaneously block diagonal, in such a way that each block represents an irreducible group of unitary operators on a smaller Hilbert space.

Reducing a unitary representation on a finite-dimensional Hilbert space into irreducible subrepresentations is a bit like decomposing an integer into a product of prime factors. As with prime factorization, the decomposition process for a finite-dimensional unitary representation has only one possible end: there is, up to ordering, a unique list

of irreducible representations into which a given unitary representation decomposes. But in infinite dimensions the decomposition process faces a number of difficulties, the most surprising of which is that there may be two decompositions of the same representation into entirely different sets of irreducible subrepresentations.

In the face of this, a different form of decomposition suggests itself, which is roughly analogous to the factorization of an integer into prime powers instead of individual primes. Let us refer to the prime powers into which an integer is decomposed as its *components*. They have two characteristic properties: no two components share a common factor, and any two (proper) factors of the same component *do* share a common factor. Similarly, one can decompose a unitary representation into *isotypical components* which have analogous properties: no two distinct isotypical components share a common (meaning ISOMORPHIC) subrepresentation, and any two subrepresentations of the same isotypical component have themselves a common sub-subrepresentation. Any unitary representation (finite-dimensional or not) can be decomposed into isotypical components, and this decomposition is unique.

In finite dimensions, every isotypical representation decomposes into a (finite) number of *identical* irreducible subrepresentations (like the prime factors of a prime power). In infinite dimensions this is not so. In effect, much of von Neumann algebra theory is concerned with analyzing the many possibilities that arise.

2.2 Factors

The commutant of an isotypical unitary representation is called a *factor*. Concretely, a factor is a von Neumann algebra M whose *center*, the set of all operators in M that commute with every member of M , consists of nothing more than scalar multiples of the identity operator. This is because projections in the center of M correspond to projections onto combinations of isotypical subrepresentations. Every von Neumann algebra can be uniquely decomposed into factors.

A factor is said to be of *type I* if it arises as the commutant of an isotypical representation which is a multiple of a single irreducible representation. Every type I factor is isomorphic to the algebra

of all bounded operators on a Hilbert space. In finite dimensions, every factor is of type I, since as we already noted every isotypical representation decomposes into a multiple of one irreducible representation.

The existence of unitary representations with more than one decomposition into irreducible components is related to the existence of factors which are *not* of type I. Von Neumann, together with Francis Murray, investigated this possibility in a series of papers that mark the foundation of operator algebra theory. They introduced an order structure on the collection of subrepresentations of a given isotypical representation or, to put it in terms of the commutant, on the collection of projections in a given factor. If H_0 and H_1 are subrepresentations of the isotypical representation H , then we write $H_0 \preceq H_1$ if H_0 is isomorphic to a subrepresentation of H_1 . Murray and von Neumann proved that this is a TOTAL ORDERING: either $H_0 \preceq H_1$; or $H_1 \preceq H_0$; or both, in which case H_0 and H_1 are isomorphic. For example, in a finite-dimensional type I situation, where H is a multiple of n copies of a single irreducible representation, each subrepresentation is the sum of $m \leq n$ copies of the irreducible representation, and the order structure of the (isomorphism classes of) subrepresentations is the same as the order structure of the integers $\{0, 1, \dots, n\}$.

Murray and von Neumann showed that the only order structures that can arise from factors are the following very simple ones:

- Type I $\{0, 1, 2, \dots, n\}$ or $\{0, 1, 2, \dots, \infty\}$;
- Type II $[0, 1]$ or $[0, \infty]$;
- Type III $\{0, \infty\}$.

The *type* of a factor is determined from the order structure of its projections according to this table.

In the case of factors of type II, the order structure is that of an interval of *real numbers*, not integers. Any subrepresentation of an isotypical representation of type II can be divided into yet smaller subrepresentations: we shall never reach an irreducible “atom”. Nevertheless, subrepresentations can still be compared in size by means of the “real-valued dimension” provided by Murray and von Neumann’s theorem.

A notable example of a factor of type II may be obtained as follows. Let G be a group and let $H = \ell^2(G)$ be a Hilbert space having basis vectors

$[g]$ corresponding to the elements $g \in G$. Then there is a natural representation of G on H derived from the group multiplication law, called the *regular representation*, which acts on the basis vectors by $U_g[g'] = [gg']$. The commutant of this representation is a von Neumann algebra M . If G is a commutative group, then all the operators U_g are in the center of M ; but if G is far enough from commutativity (for instance, if it is a free group) then M will have trivial center and will therefore be a factor. It can be shown that this factor has type II. There is a simple explicit formula for the real-valued dimension of a subrepresentation corresponding to an orthogonal projection $P \in M$. Represent P by an infinite matrix relative to the basis $\{[g]\}$ of H . Because P commutes with the representation, it is easy to see that the diagonal elements of P are all the same, equal to some real number between 0 and 1. This real number is the dimension of the subrepresentation corresponding to P .

More recently, the Murray-von Neumann dimension theory has found unexpected applications in topology. Many important topological invariants such as BETTI NUMBERS are defined as the dimensions of certain vector spaces. Using von Neumann algebras, one can define real-valued counterparts of these invariants which have useful additional properties. By relating the real-valued invariants to their ordinary counterparts, one can use von Neumann algebra theory to obtain topological conclusions. The von Neumann algebras used here are typically obtained by the construction of the previous paragraph from the FUNDAMENTAL GROUP of some compact space.

2.3 Modular theory

Type III factors remained for a long time rather mysterious; indeed, Murray and von Neumann were at first unable to determine whether any such factors existed. This they were eventually able to do, but the crucial breakthrough came well after the pioneering work of Murray and von Neumann, when it was realized that each von Neumann algebra has a special family of symmetries, its so-called *modular automorphism group*.

To explain the origins of modular theory, let us consider once again the von Neumann algebra obtained from the regular representation of a group

G . We defined the operators U_g on $\ell^2(G)$ by multiplying *on the left* by elements of G ; but we could equally well have considered a representation defined by multiplying *on the right*. This would yield a different von Neumann algebra.

So long as we deal only with discrete groups G this difference is unimportant, because the map $S: [g] \mapsto [g^{-1}]$ is a unitary operator on H which interchanges the left and right regular representations. But for certain *continuous* groups the problem arises that the function $f(g)$ may be square-integrable while $f(g^{-1})$ is not. In this situation there is no simple unitary isomorphism analogous to the one for discrete groups. To remedy this one must introduce a correction factor called the *modular function* of G .

The project of modular theory is to show that something analogous to the modular function can be constructed for any von Neumann algebra. This object then serves as an invariant for all factors of type III, whether or not they are explicitly derived from groups.

Modular theory exploits a version of the GNS construction (Section 1.4). Let M be a self-adjoint algebra of operators. A linear functional $\phi: M \rightarrow \mathbb{C}$ is called a *state* if it is positive in the sense that $\phi(T^*T) \geq 0$, for every $T \in M$ (this terminology is derived from the connection described earlier between Hilbert space theory and quantum mechanics). For the purposes of modular theory we restrict attention to *faithful* states, those for which $\phi(T^*T) = 0$ implies $T = 0$. If ϕ is a state, then the formula

$$\langle T_1, T_2 \rangle = \phi(T_1^* T_2)$$

defines an inner product on the vector space M . Applying the GNS procedure, we obtain a Hilbert space H_M . The first important fact about H_M is that every operator T in M determines an operator on H_M . Indeed a vector $V \in H_M$ is a limit $V = \lim_{n \rightarrow \infty} V_n$ of elements in M , and we can apply an operator $T \in M$ to the vector V using the formula

$$TV = \lim_{n \rightarrow \infty} TV_n,$$

where on the right hand side we use multiplication in the algebra M . Because of this observation, we can think of M as an algebra of operators on H_M , rather than as an algebra of operators on whatever Hilbert space we began with.

Next, the adjoint operation equips the Hilbert space H_M with a natural “antilinear” operator $S: H_M \rightarrow H_M$ by the formula² $S(V) = V^*$. Since $U_g^* = U_{g^{-1}}$ for the regular representation, this is indeed analogous to the operator S we encountered in our discussion of continuous groups. The important theorem of Minoru Tomita and Masamichi Takesaki asserts that, as long as the original state ϕ satisfies a continuity condition, the complex powers $U_t = (S^*S)^{it}$ have the property that $U_t M U_{-t} = M$, for all t .

The transformations of M given by the formula $T \mapsto U_t T U_{-t}$ are called the *modular automorphisms* of M . Alain Connes proved that they depend only in a rather inessential way on the original faithful state ϕ . To be precise, changing ϕ changes the modular automorphisms only by *inner automorphisms*, that is transformations of the form $T \mapsto U T U^{-1}$, where U is a unitary operator in M itself. The remarkable conclusion is that every von Neumann algebra M has a canonical one-parameter group of “outer automorphisms” which is determined by the M alone.

The modular group of a type I or type II factor consists only of the identity transformation; however, the modular group of a type III factor is much more complex. For example, the set

$$\{t \in \mathbb{R} : T \mapsto U_t T U_{-t} \text{ is an inner automorphism}\}$$

is a subgroup of \mathbb{R} and an invariant of M which may be used to distinguish among uncountably many different type III factors.

2.4 Classification

A crowning achievement of von Neumann algebra theory is the classification of factors which are *approximately finite dimensional*. These are the factors which are in some sense limits of finite-dimensional algebras. Besides the range of the dimension function, which separates factors into types, the sole invariant is the *module*, which is a flow on a space assembled from the modular automorphism group.

A lot of current attention is being given to the long-standing problem of distinguishing among the type II factors associated to the regular representations of groups. Of special interest is the case of

²The interpretation of this formula on the completion H_M of M is a delicate matter.

free groups, around which has flourished the subject of *free probability theory*. Despite intensive effort, some fundamental questions remain open: at the time of writing it is unknown whether the factors associated to the free groups on 2 and on 3 generators are isomorphic.

Another important development has been *subfactor theory* which attempts to classify the ways in which factors can be realized within other factors. A remarkable and surprising theorem of Vaughan Jones shows that in the type II situation, where continuous values of dimensions are the norm, the dimensions of subfactors can in certain situations assume only a discrete range of values. The combinatorics associated to this result have also appeared in other apparently quite unrelated parts of mathematics, notably KNOT THEORY.

3 C*-Algebras

Von Neumann algebra theory helps describe the structure of a single representation of a group on a Hilbert space. But in many situations it is of interest to gain an understanding of all possible unitary representations. To shed some light on this problem we turn to a related but different part of operator algebra theory.

Consider the collection $\mathcal{B}(H)$ of all bounded operators on a Hilbert space H . It has two very different structures: *algebraic* operations such as addition, multiplication, and formation of adjoints, and *analytic* structures such as the operator norm

$$\|T\| = \sup\{\|Tu\| : \|u\| \leq 1\}.$$

These structures are not independent of one another. Suppose for instance that $\|T\| < 1$ (an analytic hypothesis). Then the geometric series

$$S = I + T + T^2 + T^3 + \dots$$

converges in $\mathcal{B}(H)$, and its limit S satisfies

$$S(I - T) = (I - T)S = I.$$

It follows that $I - T$ is invertible in $\mathcal{B}(H)$ (an algebraic conclusion). One can easily deduce from this that the *spectral radius* $r(T)$ of any operator T (defined to be the greatest absolute value of any complex number in the spectrum of T) is less than or equal to its norm.

The remarkable *spectral radius formula* goes much further in the same direction. It asserts that $r(T) = \lim_{n \rightarrow \infty} \|T^n\|^{1/n}$. If T is *normal* ($TT^* = T^*T$), and in particular if T is self-adjoint, then it may be shown that $\|T^n\| = \|T\|^n$. As a result, the spectral radius of T is precisely equal to the norm of T . There is therefore a very close connection between the algebraic structure of $\mathcal{B}(H)$, particularly algebraic structure related to the adjoint operation, and the analytic structure.

Not all the properties of $\mathcal{B}(H)$ are relevant to this connection between algebra and analysis. A *C*-algebra* A is an abstract structure which has enough properties to allow the argument of the previous two paragraphs to go through. A detailed definition would be out of place here, but it is worth mentioning that a crucial condition relating norm, multiplication and *-operation is

$$\|a^*a\| = \|a\|^2, \quad a \in A,$$

called the *C*-identity* for A . We note also that special classes of operators on Hilbert space (unitaries, orthogonal projections, and so on) all have their counterparts in a general C*-algebra. For example, a *unitary* $u \in A$ satisfies $uu^* = u^*u = 1$, and a *projection* p satisfies $p = p^2 = p^*$.

A simple example of a C*-algebra is obtained by starting with a single operator $T \in \mathcal{B}(H)$. The collection of all operators $S \in \mathcal{B}(H)$ that can be obtained as limits of polynomials in T and T^* is a C*-algebra said to be *generated* by T . The C*-algebra generated by T is commutative if and only if T is normal; this is one reason for the importance of normal operators.

3.1 Commutative C*-algebras

If X is a compact topological space, then the collection $C(X)$ of continuous functions $f: X \rightarrow \mathbb{C}$ has algebraic operations (inherited from the usual ones on \mathbb{C}) and a norm $\|f\| = \sup\{|f(x)| : x \in X\}$. In fact, these operations make $C(X)$ into a C*-algebra. The multiplication in $C(X)$ is *commutative*, because the multiplication of complex numbers is commutative.

A basic result of Gelfand and Naimark asserts that *every* commutative C*-algebra is isomorphic to some $C(X)$. Given a commutative C*-algebra A , one constructs X as the collection of all algebra homomorphisms $\xi: A \rightarrow \mathbb{C}$, and the *Gelfand*

transform then associates to $a \in A$ the function $\xi \mapsto \xi(a)$ from X to \mathbb{C} .

The Gelfand-Naimark theorem is a foundational result of operator theory. For example, a modern proof of the spectral theorem might proceed as follows. Let T be a self-adjoint or normal operator on a Hilbert space H , and let A be the commutative C*-algebra generated by T . By the Gelfand-Naimark theorem A is isomorphic to $C(X)$ for some space X , which may in fact be identified with the spectrum of T . If v is a unit vector in H , then the formula $S \mapsto \langle Sv, v \rangle$ defines a state ϕ on A . The GNS space associated to this state is a Hilbert space of functions on X , and elements of $A = C(X)$ act as multiplication operators. In particular, T acts as a multiplication operator. A small additional argument shows that T is unitarily equivalent to this multiplication operator, or at least to a direct sum of such operators (which is itself a multiplication operator on a larger space).

Continuous functions can be composed: if f and g are continuous functions (with the range of g contained in the domain of f) then $f \circ g$ is also a continuous function. Since the Gelfand-Naimark theorem tells us that any self-adjoint element of a C*-algebra A sits inside an algebra isomorphic to the continuous functions on the spectrum of a , we conclude that if $a \in A$ is self-adjoint, and if f is a continuous function on the spectrum of a , then an operator $f(a)$ exists in A . This *functional calculus* is a key technical tool in C*-algebra theory. For example, suppose that $u \in A$ is unitary and $\|u - 1\| < 2$. Then the spectrum of u is a subset of the unit circle in \mathbb{C} that does not contain -1 . One can define a continuous branch of the complex logarithm function on such a subset, and it follows that there is $a = \log u$, $a = -a^*$, such that $u = e^a$. The path $t \mapsto e^{ta}$, $0 \leq t \leq 1$, is then a continuous path of unitaries in A connecting u to the identity. Thus every unitary sufficiently close to the identity is connected to the identity by a unitary path.

3.2 Further examples of C*-algebras

3.2.1 The compact operators

An operator on a Hilbert space has *finite rank* if its range is a finite-dimensional subspace. The operators of finite rank form an algebra, and its closure is a C*-algebra called the algebra of *compact* op-

erators and denoted \mathcal{K} . One can also view \mathcal{K} as a “limit” of matrix algebras

$$M_1(\mathbb{C}) \rightarrow M_2(\mathbb{C}) \rightarrow M_3(\mathbb{C}) \rightarrow \dots$$

where each matrix algebra is included in the next by

$$A \mapsto \begin{pmatrix} A & 0 \\ 0 & 0 \end{pmatrix}.$$

Many natural operators are compact, including the integral operators that arose in Fredholm’s theory. The identity operator on a Hilbert space is compact if and only if that Hilbert space is finite-dimensional.

3.2.2 The CAR algebra

The presentation of \mathcal{K} as a limit of matrix algebras leads one to consider other “limits” of a similar sort. (We shall not attempt a formal definition of these limits here, but it is important to note that the limit of a sequence $A_1 \rightarrow A_2 \rightarrow A_3 \rightarrow \dots$ depends on the homomorphisms $A_i \rightarrow A_{i+1}$ as well as on the algebras A_i .) One particularly important example is obtained as the limit

$$M_1(\mathbb{C}) \rightarrow M_2(\mathbb{C}) \rightarrow M_4(\mathbb{C}) \rightarrow \dots$$

where each matrix algebra is included in the next by

$$A \mapsto \begin{pmatrix} A & 0 \\ 0 & A \end{pmatrix}.$$

This is called the *CAR algebra*, because it contains elements that represent the *canonical anti-commutation relations* which arise in quantum theory. C^* -algebras find several applications to quantum field theory and quantum statistical mechanics which extend von Neumann’s formulation of quantum theory in terms of Hilbert space.

3.2.3 Group C^* -algebras

If G is a group and $g \mapsto U_g$ is a unitary representation of G on a Hilbert space H , we can consider the smallest C^* -algebra of operators on H containing all the U_g ; this is called the C^* -algebra *generated* by the representation. An important example is the *regular representation* on the Hilbert

space $\ell^2(G)$ generated by G , which we defined in Section 2.2. The C^* -algebra that it generates is denoted $C_r^*(G)$. The subscript r refers to the regular representation. Considering other representations leads to other, potentially different, group C^* -algebras.

Consider for example the case $G = \mathbb{Z}$. Since this is a commutative group, its C^* -algebra is also commutative, and thus it is isomorphic to $C(X)$ for a suitable X , by the Gelfand-Naimark theorem. In fact, X is the unit circle S^1 , and the isomorphism

$$C(S^1) \cong C_r^*(\mathbb{Z})$$

takes a function on the circle to its Fourier series.

States on group C^* -algebras correspond to positive-definite functions on groups, and hence to unitary group representations. In this way new representations may be constructed and studied. For example, using states of group C^* -algebras it is possible to give the set of irreducible representations of G the structure of a topological space.

3.2.4 The irrational rotation algebra

The algebra $C^*(\mathbb{Z})$ is generated by a single unitary element U (corresponding to $1 \in \mathbb{Z}$). Moreover, it is the *universal example* of such a C^* -algebra, which is to say that given any C^* -algebra A and unitary $u \in A$, there is one and only one homomorphism $C^*(\mathbb{Z}) \rightarrow A$ sending U to u . In fact, this is nothing other than the functional calculus homomorphism for the unitary u .

If instead we consider the universal example of a C^* -algebra generated by *two* unitaries U, V subject to the relation

$$UV = e^{2\pi i\alpha} VU,$$

where α is irrational, we obtain a noncommutative C^* -algebra called the *irrational rotation algebra* A_α . The irrational rotation algebras have been studied intensively from a number of points of view. Using K -theory (see below) it has been shown that A_{α_1} is isomorphic to A_{α_2} if and only if $\alpha_1 \pm \alpha_2$ is an integer.

It can be shown that the irrational rotation algebra is *simple*, which implies that *any* pair of unitaries U, V satisfying the commutation relation above generate a copy of A_α . (Note the strong

contrast to the case of a single unitary: 1 is a unitary operator, but it does not generate a copy of $C^*(\mathbb{Z})$.) This allows us to give a concrete representation of A_α on the Hilbert space $L^2(S^1)$, where U is the rotation through $2\pi\alpha$ and V is multiplication by $z: S^1 \rightarrow \mathbb{C}$.

4 Fredholm operators

A *Fredholm operator* between Hilbert spaces is a bounded operator T for which the kernel and cokernel are finite-dimensional. This means that the homogeneous equation $Tu = 0$ admits only finitely many linearly independent solutions, while the inhomogeneous equation $Tu = v$ admits a solution if v satisfies a finite number of linear conditions. The terminology arises from Fredholm's original work on integral equations; he showed that if K is an integral operator, then $I + K$ is a Fredholm operator.

For the operators that Fredholm considered, the dimensions of the kernel and cokernel must be equal, but in general this need not be so. The *unilateral shift operator* S which maps the infinite "row vector" (a_1, a_2, a_3, \dots) to $(0, a_1, a_2, \dots)$ is an example. The equation $Su = 0$ has only the zero solution, but the equation $Su = v$ has a solution only if the first coordinate of the vector v is zero.

The *index* of a Fredholm operator is defined to be the integer difference

$$\text{Index}(T) = \dim(\ker(T)) - \dim(\text{coker}(T)).$$

For example, every invertible operator is a Fredholm operator of index 0, whereas the unilateral shift is a Fredholm operator of index -1 .

4.1 Atkinson's theorem

Consider the two systems of linear equations

$$\begin{array}{l} 2.1x + y = 0 \\ 4x + 2y = 0 \end{array} \quad \text{and} \quad \begin{array}{l} 2x + y = 0 \\ 4x + 2y = 0 \end{array} .$$

Although the coefficients of these equations are very close, the dimensions of their kernels are quite different: the left-hand system has only the zero solution, whereas the right-hand system has the non-trivial solutions $(t, -2t)$. Thus the dimension of the kernel is an *unstable* invariant of the system of equations. A similar remark applies to the

dimension of the cokernel. By contrast the index is stable, despite its definition as the difference of two unstable quantities.

An important theorem of Frederick Atkinson gives precise expression to these stability properties. Atkinson's theorem asserts an operator T is Fredholm if and only if it is invertible modulo compact operators. This implies that any operator that is sufficiently close to a Fredholm operator is itself a Fredholm operator with the same index, and that if T is a Fredholm operator and K is a compact operator, then $T + K$ is a Fredholm operator with the same index as T . Notice that, since integral operators are compact operators, this contains Fredholm's original theorem as a special case.

4.2 The Toeplitz index theorem

TOPOLOGY studies those properties of mathematical systems that remain the same when the system is (continuously) perturbed. Atkinson's theorem tells us that the Fredholm index is a topological quantity. In many contexts it is possible to obtain a formula for the index of a Fredholm operator in terms of other, apparently quite different, topological quantities. Formulas of this sort often indicate deep connections between analysis and topology and often have powerful applications.

The simplest example involves the *Toeplitz operators*. A Toeplitz operator has a matrix with the special form

$$T = \begin{pmatrix} b_0 & b_1 & b_2 & b_3 & \cdots \\ b_{-1} & b_0 & b_1 & b_2 & \cdots \\ b_{-2} & b_{-1} & b_0 & b_1 & \cdots \\ b_{-3} & b_{-2} & b_{-1} & b_0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} .$$

The entries are constant down each diagonal. The sequence of coefficients $\{b_n\}_{n=-\infty}^{\infty}$ defines a function $f(z) = \sum_{n=-\infty}^{\infty} b_n z^{-n}$ on the unit circle in the complex plane, called the *symbol* of the Toeplitz operator. It can be shown that a Toeplitz operator whose symbol is a continuous function which is never zero is Fredholm. What is its index?

The answer is given by thinking about the symbol as a mapping from the unit circle to the nonzero complex numbers—in other words, as a closed path in the nonzero complex plane. The fundamental topological invariant of such a path is

its *winding number*—the number of times it “goes around” the origin in the counterclockwise direction. It can be proved that the index of a Toeplitz operator with nonzero symbol f is minus the winding number of f . For example, if f is the function $f(z) = z$ (with winding number $+1$), then the associated Toeplitz operator is the unilateral shift S that we encountered earlier (with index -1). The Toeplitz index theorem is a very special case of the ATIYAH-SINGER INDEX THEOREM which gives a topological formula for the indices of various Fredholm operators arising from geometry.

4.3 Essentially normal operators

Atkinson’s theorem suggests that compact perturbations of an operator are in some sense “small”. This leads to the study of properties of an operator that are preserved by compact perturbation. For instance, the *essential spectrum* of an operator T is the set of complex numbers λ for which $T - \lambda I$ fails to be Fredholm (that is, invertible modulo compact operators). Two operators T_1 and T_2 are *essentially equivalent* if there is a unitary operator U such that UT_1U^* and T_2 differ by a compact operator. A beautiful theorem originally due to Hermann Weyl then asserts that two self-adjoint or normal operators are essentially equivalent if and only if they have the same essential spectrum.

One might argue that the restriction to normal operators in this theorem is inappropriate. Since we are concerned with properties that are preserved by compact perturbation, would it not be more appropriate to consider *essentially normal* operators—operators T for which $T^*T - T^*T$ is compact? This apparently modest variation leads to an unexpected result. The unilateral shift S is an example of an essentially normal operator. Its essential spectrum is the unit circle, as is the essential spectrum of its adjoint; however, S and S^* cannot be essentially equivalent, because S has index -1 and S^* has index $+1$. Thus some new ingredient, beyond the essential spectrum, is needed to classify essentially normal operators. In fact, it follows easily from Atkinson’s theorem that if essentially normal operators T_1 and T_2 are to be essentially equivalent, then not only must they have the same essential spectrum but also, for every λ not in the essential spectrum, the Fredholm index of $T_1 - \lambda I$ must equal to the Fredholm index of

$T_2 - \lambda I$. The converse to this statement was proved by Larry Brown, Ron Douglas and Peter Fillmore in the 1970s, using entirely novel techniques which led to a new era of interaction between C^* -algebra theory and topology.

4.4 K-Theory

A remarkable feature of the Brown-Douglas-Fillmore work was the appearance within it of tools from algebraic topology, notably *K-theory*. Remember that according to the Gelfand-Naimark theorem, the study of (suitable) topological spaces and the study of commutative C^* -algebras are one and the same; all the techniques of topology can be transferred, via the Gelfand-Naimark isomorphism, to commutative C^* -algebras. Having made this observation, it is natural to ask which of these techniques can be extended further, to provide information about *all* C^* -algebras, commutative or not. The first and best example is *K-theory*.

In its most basic form, *K-theory* associates to each C^* -algebra A an abelian group $K(A)$, and to each homomorphism of C^* -algebras a corresponding homomorphism of abelian groups. The building blocks for $K(A)$ can be thought of as generalized Fredholm operators associated to A ; the generalization is that these operators act on Hilbert “spaces” in which the complex scalars are replaced by elements of the C^* -algebra A . The group $K(A)$ itself is defined to be the collection of connected components of the space of all such generalized Fredholm operators. Thus if $A = \mathbb{C}$, for instance (so that we are dealing with classical Fredholm operators), then $K(A) = \mathbb{Z}$. This follows from the fact that two Fredholm operators are connected by a path of Fredholm operators if and only if they have the same index.

One of the great strengths of *K-theory* is that *K-theory* classes may be constructed from a variety of different ingredients. For example, every projection $p \in A$ defines a class in $K(A)$ which can be thought of as a “dimension” for the range of p . This connects *K-theory* to the classification of factors (Section 2.2), and has become an important tool in the effort to classify various families of C^* -algebras—such as the irrational rotation algebras. (It was at one time thought that the irrational rotation algebras might not contain any nontrivial projections at all: the construction of such projec-

tions by Marc Rieffel was an important step in the development of C^* -algebra K -theory.) Another beautiful example is George Elliott’s classification theorem for locally finite dimensional C^* -algebras like the CAR algebra—they are completely determined by K -theoretic invariants.

The problem of computing the K -theory groups of noncommutative C^* -algebras, particularly group C^* -algebras, has turned out to have important connections with topology. In fact some key advances in topology have come from C^* -algebra theory in this way, thereby allowing operator algebraists to repay some of the debt they owe to the topologists for K -theory. The principal organizing problem in this area is the *Baum-Connes conjecture*, which proposes a description of the K -theory of group C^* -algebras in terms of invariants familiar in algebraic topology. Most of the progress on the conjecture to date is the result of work of Gennadi Kasparov, who dramatically broadened the original discoveries of Brown, Douglas, and Fillmore to cover not just single essentially normal operators but noncommuting systems of operators, that is, C^* -algebras. Kasparov’s work is now a central component of operator algebra theory.

5 Noncommutative geometry

Descartes’ invention of coordinates showed that one can do geometry not by thinking directly about the points of space and their interrelationships, but instead by thinking about coordinate functions defined on space: these functions being, of course, the familiar x , y and z . The Gelfand-Naimark theorem can be viewed as one expression of this idea of passing from the “point picture” of a space X to the “field picture” of the algebra $C(X)$ of functions on it. The success of K -theory in operator algebras invites us to ponder whether the field picture might be *more powerful* than the point picture, since K -theory can be applied to noncommutative C^* -algebras which may not have any “points” (homomorphisms to \mathbb{C}) at all.

One of the most exciting research frontiers in operator algebra theory is reached along a path which develops these thoughts. The *noncommutative geometry* program of Alain Connes takes seriously the idea that a general C^* -algebra should be

thought of as an algebra of functions on a “non-commutative space”, and goes on to develop “non-commutative” versions of many ideas from geometry and topology, as well as completely new constructions that have no commutative counterpart. Noncommutative geometry begins with the creative reformulation of ideas from ordinary geometry in ways which involve only operators and functions, but not points.

Consider for instance the circle S^1 . The algebra $C(S^1)$ reflects all the topological properties of S^1 , but, to distinguish its *metric* (distance-related) properties also, we look at not just $C(S^1)$ but the pair consisting of the algebra $C(S^1)$ and the operator $D = id/d\theta$ on the Hilbert space $H = L^2(S^1)$. Notice that if f is a function on the circle (considered as a multiplication operator on H), then the commutator $Df - fD$ is also a multiplication operator, this time by $idf/d\theta$. It follows that ordinary measurements of angular distance between points on the circle can be recovered from $C(S^1)$ and D by the formula

$$d(p, q) = \max\{|f(p) - f(q)| : \|Df - fD\| \leq 1\}.$$

Connes argues that operator $|D|^{-1}$ plays the role of the “unit of arc-length ds ” in this and many other more complicated situations.³

Another feature of the examples Connes considers, also of central importance in noncommutative geometry, is the fact that the $|D|^{-k}$ is a trace-class operator (see Section 1.5) when k is large enough—in the case of the circle, when k is bigger than 1. Computations with traces connect noncommutative geometry to COHOMOLOGY THEORY. We now have two kinds of “noncommutative algebraic topology”, namely K -theory and a new variant of homology called *cyclic cohomology*; the connection between the two is provided by a very general index theorem.

There are several procedures which produce noncommutative C^* -algebras (to which Connes’ methods can be applied) from classical geometric data. The irrational rotation algebras A_θ are examples; the classical picture to which they apply is the QUOTIENT SPACE of the circle by the group of rotations through multiples of θ . Classical methods

³The operator D not quite invertible since D vanishes on constant functions. A small modification must therefore be made before considering inverse operators. The operator $|D|$ is by definition the positive square root of D^2

of geometry and topology are unable to handle this quotient space, but the noncommutative approach via A_θ is much more successful.

An exciting but speculative possibility is that the basic laws of physics should be addressed from the perspective of noncommutative geometry. The transition to noncommutative C^* -algebras can be viewed as analogous to the transition from classical to quantum mechanics. However, Connes has argued that noncommutative C^* -algebras play a role in describing the physical world even before the transition is made to quantum physics.

Bibliography

1. Connes, A. 1995, *Non-commutative geometry*. Academic Press, Boston.
2. Davidson, K. 1996, *C^* -algebras by example*. American Mathematical Society, Providence, RI.
3. Fillmore, P. 1996, *A user's guide to operator algebras*. Canadian Mathematical Society Series of Monographs and Advanced Texts, John Wiley & Sons Inc., New York.
4. Halmos, P. R. 1963, *What does the spectral theorem say?*, Amer. Math. Monthly **70**, 241–247

Biography of contributors

Nigel Higson and John Roe are professors of mathematics at Penn State.