Analytic K-Homology

Nigel Higson and John Roe

Introduction

Analytic K-homology draws together ideas from algebraic topology, functional analysis and geometry. It is a tool — a means of conveying information among these three subjects — and it has been used with spectacular success to prove and indeed discover remarkable theorems across a wide span of mathematics. These include results in operator theory which make no mention of topology or geometry at all, and results in topology and geometry which are apparently far removed from functional analysis.

The purpose of this book is to acquaint the reader with the essential ideas of analytic K-homology and develop some of its applications. We shall begin this introduction with a brief account of the origins of the subject, followed by a first look at the central notion of 'abstract elliptic operator' which ties most of the different aspects of analytic K-homology together.

The roots of analytic K-homology

The subject of analytic K-homology had two separate beginnings, one in operator theory and one in the index theory of Atiyah and Singer, but we shall start this story in a third place, in the realm of algebraic topology.

The abstract machinery of algebraic topology assigns to each generalized *co-homology* theory a dual *homology* theory. From this point of view, K-homology is nothing more than the homology theory dual to Atiyah–Hirzebruch K-theory. Now cohomology and homology theories stand in relation to one another more or less as vector spaces stand in relation to their duals. Thus an essential feature of any dual pair of homology/cohomology theories is the bilinear pairing between the homology and cohomology groups of each space, the exemplar of which is the pairing between ordinary homology and de Rham cohomology given by integration of differential forms over simplices. An important step in the investigation of any homology/cohomology theory is to find a concrete realization of this pairing, just as the analysis of a linear space is often centered (especially in the case of topological linear spaces) around a concrete identification of its dual.

Unfortunately the abstract approach to dual theories does not provide a concrete and geometric account of this pairing. The first purpose of analytic K-homology is to remedy this shortcoming in the case of K-theory: at its core is a very simple pairing based on Fredholm index theory. But what makes analytic K-homology especially remarkable, and much more than an appendix to Atiyah–Hirzebruch K-theory, is the extraordinary volume of mathematics which travels in both directions between topology and functional analysis through this basic Fredholm index pairing.

The Fredholm index of a linear operator T is the integer quantity

$$\text{Index}(T) = \text{Dim}(\text{Kernel}(T)) - \text{Dim}(\text{Cokernel}(T)),$$

which is defined whenever the kernel and cokernel of T are finite-dimensional. The index is stable under finite-rank perturbations of the operator T (in finite dimensions this is the well-known Rank-Nullity Theorem from college linear algebra). In the context of Hilbert space, where this book is located, the index is also stable under small-norm perturbations of the bounded operator T. Since every compact Hilbert space operator is a norm-limit of finite rank operators, it follows that the Fredholm index is stable under compact operator perturbations of T. In fact a lovely theorem of Atkinson asserts that a bounded Hilbert space operator has finite-dimensional kernel and cokernel if and only if it is invertible modulo compact operators. (The connection between index theory and compact operators is already made evident by the famous Fredholm Alternative: if K is a compact operator then the linear equation $Kf + f = g$ has a solution for every g if and only if the homogeneous equation $Kf + f = 0$ has no non-trivial solution.)

Operator theory has long considered the problem of classifying Hilbert space operators 'modulo compact operators'. Weyl and von Neumann showed that two selfadjoint operators are unitarily equivalent modulo compact operators if and only if they have the same spectrum apart from isolated eigenvalues of finite multiplicity. In the 1960's Brown, Douglas and Fillmore began an investigation of *essentially normal* operators, meaning those for which T^*T and TT^* are equal modulo compact operators, by asking themselves the following question: is the unilateral shift operator on the Hilbert space $\ell^2(\mathbb{N})$ unitarily equivalent modulo compact operators to the bilateral shift operator on $\ell^2(\mathbb{Z})$? The essential spectrum, meaning the part of the spectrum which is stable under compact perturbations, is for both operators the unit circle S^1 in the complex plane. According to the Weyl–von Neumann Theorem the essential spectrum is a complete classification invariant for selfadjoint operators. But in the present case a new invariant emerges, namely the Fredholm index. Indeed the index of the unilateral shift is -1, whereas the index of the bilateral shift is 0, while the stability properties of the index show that it is an invariant for unitary equivalence modulo compact operators.

Using simple operator theory techniques it is not hard to show that two essentially normal operators with essential spectrum S^1 are unitarily equivalent modulo compact operators if and only if they have the same Fredholm index. But the situation for other essential spectra $X \subseteq \mathbb{C}$ (for example, the closed unit disk) is considerably more complicated. Brown, Douglas and Fillmore introduced the classifying structure $\mathrm{Ext}(X)$ to help attack the problem, and then they proved two very unexpected things: first, $\mathrm{Ext}(X)$ is actually an abelian group, and second, $\mathrm{Ext}(X)$ is the degree-one K-homology group of X. This is the Brown–Douglas–Fillmore Theorem. The determination of $\mathrm{Ext}(X)$, which is to say the classification of essentially normal operators, was thereby carried out by reducing the classification problem to a computational problem in algebraic topology.

The group $\mathrm{Ext}(X)$ also classifies C^*-algebra extensions of the form

$$0 \longrightarrow \mathcal{K}(H) \longrightarrow E \longrightarrow C(X) \longrightarrow 0,$$

where $\mathcal{K}(H)$ is the C^*-algebra of compact operators on a Hilbert space H and $C(X)$ is the commutative C^*-algebra of continuous complex-valued functions on X (which can now be any compact metric space). Thinking of $\mathrm{Ext}\,X)$ from this point of view, we can give a very simple account of the crucial Fredholm index pairing with K-theory. Indeed the K-theory group $K^{-1}(X)$ is generated by homotopy classes of maps from X to the group of invertible complex matrices. Given an extension, as above, an invertible, matrix-valued function on X lifts to a matrix over the algebra E which is invertible modulo compact operators. This matrix has a Fredholm index, and thanks to the stability properties of the index this procedure for combining an extension and an invertible, matrix-valued function defines a bilinear pairing between $\mathrm{Ext}(X)$ and $K^{-1}(X)$. From this point of view, the Brown–Douglas–Fillmore Theorem is the assertion that if $X \subseteq \mathbb{C}$ then the homomorphism

$$\mathrm{Index}\colon\ \mathrm{Ext}(X) \longrightarrow \mathrm{Hom}(K^{-1}(X), \mathbb{Z})$$

is an isomorphism of abelian groups.

The index theory of Atiyah and Singer presents a second view of the Fredholm index pairing between K-theory and K-homology. Suppose that X is a compact manifold and that D is an linear elliptic operator on X. Then D has a Fredholm index. But in addition if V is a vector bundle on X then a standard construction in index theory (essentially a tensor product) produces a new linear elliptic operator D_V, 'with coefficients in V', and the assignment $V \mapsto \operatorname{Index} D_V$ determines a homomorphism

$$\operatorname{Index}_D \colon K^0(X) \longrightarrow \mathbb{Z}$$

In order to extend this discussion to spaces other than manifolds, Atiyah identified the key functional-analytic properties of an elliptic operator on a manifold and so developed an abstract notion of elliptic operator, now called a *Fredholm module*. Kasparov developed Atiyah's idea and showed that the abelian group generated by homotopy classes of Fredholm modules is another analytic model for K-homology, this time for the degree-zero K-homology group of X.

Kasparov's K-homology has proved to be an extremely powerful and flexible tool in index theory. For example the proof of the Atiyah–Singer Index Theorem itself can be presented very simply and conceptually using the product structure on K-homology. Moreover Kasparov's work has allowed a considerable strengthening of the index theory of Atiyah and Singer. Kasparov developed his theory as a tool in differential topology, and indeed some of the most powerful theorems in the topological theory of manifolds (pertaining particularly to the Novikov conjecture) rely very heavily on Kasparov's machinery. In several cases no proofs of these theorems are known which do not employ functional analysis to a very considerable extent. Thus, thanks to Kasparov's discoveries, functional analysis has repaid to topology the debt incurred by Brown, Douglas and Fillmore!

Abstract Elliptic Operators

Let us describe in a little more detail the Fredholm modules that Atiyah invented. In an effort to make the main idea as plain as possible we shall take a look at Fredholm modules which are associated not to a space X but to a group G. One of the virtues of the functional-analytic approach is that these two very different types of objects — spaces and groups — can be placed on more or less the same footing. Both give rise to C^*-algebras: the commutative C^*-algebras $C(X)$ in the first case and the group C^*-algebras $C^*(G)$ in the second.

The *unitary representation ring* $R(G)$ of G is the abelian group whose generators are finite-dimensional unitary representations V of G, with the relations

$$[V \oplus W] = [V] + [W].$$

An element of $R(G)$ can always be represented by a formal difference $[V_0] - [V_1]$ of finite-dimensional unitary representations. Note that if we were to omit the finite-dimensionality restriction then $R(G)$ would collapse to zero, thanks to the calculation

$$[V] + [V_\infty] = [V \oplus V_\infty] = [V_\infty],$$

where $V_\infty = \bigoplus^\infty V$. Nevertheless, one *can* make sense of a formal difference $[V_0] - [V_1]$ of infinite-dimensional representations of G, if along with the two representations V_0 and V_1 themselves there is given an 'approximate isomorphism' between them. This is what a Fredholm module provides. To be precise, a *Fredholm module* over G consists of a pair of Hilbert space representations of G, V_0 and V_1, together with an operator $U\colon V_0 \to V_1$ which is unitary modulo compact operators and also an intertwiner modulo compact operators (meaning that $Ug - gU$ is compact for each $g \in G$). From these objects one can construct a *Fredholm representation ring*[1] $R_{\text{Fred}}(G)$ in more or less the same way that $R(G)$ is constructed. It is plain that there is a natural map $R(G) \to R_{\text{Fred}}(G)$, but it is typically far from an isomorphism (although it *is* an isomorphism if G is finite or compact).

Kasparov's group $K_0(X)$ is defined using Fredholm modules over X, which are pairs (V_0, V_1) of Hilbert space representations of the C^*-algebra $C(X)$ equipped with a linear operator $U\colon V_0 \to V_1$ which is an 'approximate isomorphism' in the same sense as above. What Atiyah observed is that each order zero elliptic pseudodifferential operator on a closed manifold X gives rise to a Fredholm module, and so Kasparov's theory in some sense incorporates the index theory of elliptic operators on X. What became evident after the Brown–Douglas–Fillmore work came to the attention of Atiyah and Singer is that each C^*-algebra extension of the type we considered earlier corresponds to a *selfadjoint* Fredholm module — which is an approximate self-symmetry of a *single* Hilbert space representation. So in retrospect the single notion of Fredholm module very elegantly connects together the two realizations of analytic K-homology.

[1] Both $R(G)$ and $R_{\text{Fred}}(G)$ are commutative rings; the former by tensor product of representations and the latter by a much more elaborate construction — the so-called Kasparov product.

About this book

We turn now to a description of this book. The first part (Chapters 1–7) leads toward a proof of the Brown–Douglas–Fillmore Theorem. Chapter 1 introduces the basic concepts of C^*-algebra theory and operator theory. In Chapter 2 we develop the theory of the Fredholm index, which, as we have indicated, is fundamental to analytic K-homology. The chapter also provides the basic definitions of C^*-algebra extension theory. Chapter 3 is devoted to a number of technical results in C^*-algebra theory: Stinespring's Theorem, nuclearity and Voiculescu's Theorem are some of the topics covered here. The fourth chapter is devoted to K-theory for C^*-algebras. The pace here is brisk, as many readers will already be familiar with K-theory, but the chapter is self-contained and could also serve as a rapid introduction to the subject.

In Chapter 5 we begin the study of K-homology proper. We define the K-homology of a C^*-algebra A in terms of the K-theory of a suitable dual algebra $\mathcal{D}(A)$. We use the technical results of Chapter 3 to develop the analysis of these dual algebras, and we obtain the long exact sequence of K-homology associated to a C^*-algebra extension. Chapter 6 makes the connection between K-homology and *coarse geometry* — the study of geometric spaces from the point of view of their large-scale structure. The connection was originally explored from the perspective of index theory on non-compact manifolds — we shall return to this subject in the second half of the book — but our immediate business is to use coarse geometry to prove the *homotopy invariance* of K-homology, and thereby show that K-homology is a generalized homology theory in the sense of algebraic topology. By this stage we shall have assembled nearly all the tools we require to prove the Brown–Douglas–Fillmore Theorem. In Chapter 7 we complete the proof, assuming a certain naturality property of the index pairing between K-theory and K-homology. This naturality property is given a direct proof in Chapter 8, but for the reader who wants to see the BDF Theorem proved without getting involved in the second half of the book, we indicate a more circuitous way around the problem in the exercises to Chapter 7. We also prove the Universal Coefficient Theorem for K-homology, which can be thought of as a generalization of the BDF Theorem to higher-dimensional spaces X.

The second half of the book is centered around index theory. In Chapter 8 we introduce Kasparov's definition of K-homology in terms of Fredholm modules, and we show that his definition is equivalent to the duality-based one of Chapter 5. Then we carry out some key computations involving boundary maps and the index pairing. In Chapter 9 we describe the product structure on K-homology.

This complicated but powerful construction is Kasparov's major contribution to the theory: it allows very simple proofs of the main properties of K-homology, and as we shall see it also connects in a very beautiful way with the theory of elliptic operators. That theory is the subject of Chapter 10; after reviewing the basic results of elliptic operator theory on manifolds, we show how an elliptic operator gives rise to a Fredholm module and therefore to a K-homology class, and how the Kasparov product on K-homology corresponds to the external product defined by Atiyah on the class of elliptic operators. In Chapter 11 we apply K-homology to prove the Atiyah–Singer Index Theorem, at least in an illustrative special case. The proof could be put onto one line — an indication of the power of K-homology theory. We also use K-homology to prove some related index theorems: the Toeplitz Index Theorem for strongly pseudoconvex domains in $\mathbb{C}^n$, and (in an exercise) the Callias Index Theorem for operators of 'Dirac–Schrödinger' type. All these results are rather simple consequences of basic calculations of K-homology theory.

Finally, in Chapter 12, we introduce the topic of higher index theory. Here one contemplates an 'index' which is no longer an integer but an element of a C^*-algebra K-theory group. Higher index theory turns out to be critically important to a number of geometric problems, of which we have selected the positive scalar curvature problem — which manifolds carry positive scalar curvature metrics? — as an illustrative and important example. Central to higher index theory are the *Baum–Connes conjectures*. Using our work on coarse geometry we shall be able to state the conjectures, prove them in some special cases, and describe their relationship to the positive scalar curvature problem.

Some topics which are not covered in the book

The plan of this book changed several times over several years of writing. But right from the beginning we decided to avoid any serious discussion of Kasparov's bivariant KK-theory. We did so because we wanted to minimize the C^*-algebra background demanded of the reader, and because we wanted to emphasize, at least at the beginning of the book, the Brown–Douglas–Fillmore theory, which is not really illuminated by Kasparov's two-variable theory. But by avoiding KK-theory we have occasionally been forced to develop circuitous arguments at points where the bivariant theory provides a direct approach. In addition, KK-theory is more or less essential for the further development of the ideas introduced in Chapter 12. Other topics, such as the foliation index theorem of Connes and Skandalis, are also out of convenient reach.

In recent years other definitions of analytic K-homology have been developed. Perhaps most notable among them is a version of K-homology for C^*-algebras which is based on a notion of *asymptotic morphism* between C^*-algebras. While this version of K-homology has an appealing simplicity, and while it is quite well adapted to Atiyah–Singer index theory, it is less well suited to the Brown–Douglas–Fillmore theory. Moreover the theory of asymptotic morphisms is best developed in the context of a two-variable theory, like Kasparov's KK-theory, and since we are concentrating on the one-variable theory it did not seem altogether appropriate to include an account of the theory of asymptotic morphisms in this book.

After some hesitation, we decided not to include any extended discussion of *real* C^*-algebras and their K-theory and K-homology groups in the main text of the book. The connections between K-homology and manifold theory attain their most precise form if one works in the real case; but the underlying linear algebra is more complicated, and in places the more versatile and powerful bivariant KK-theory of Kasparov is called for. This means that the real case is perhaps more suited to a second look at the subject. We have however attempted to write a book which, while not real, can be 'realized' without excessive difficulty. Following a well-known precedent, we have included in Appendix B an outline of how this is to be achieved.

Advice to the reader

To read the first part of this book you will need to have completed a basic course in functional analysis, and have had some exposure to 'homological' mathematics. The latter might come from studying classical algebraic topology, or homological algebra, or the K-theory of C^*-algebras. Chapters 1–7 then constitute a somewhat discursive introduction to K-homology, culminating in the proof of the Brown–Douglas–Fillmore Theorem. The pace is not too hurried and the arguments are mostly rather detailed. These chapters could be the core of an introductory graduate course on K-homology.

The pace quickens in the second half of the book with the introduction of Kasparov theory, and from Chapter 8 on a bit more mathematical sophistication is demanded of you. For example, in many arguments you will need to supply at least some of the details yourself. Moreover, since the central examples of Fredholm modules come from elliptic operator theory, some familiarity with the machinery of smooth manifolds, vector bundles, and so on will be required. Finally several sections require an acquaintance with specialized topics in Riemannian geometry, complex function theory, and other subjects.

Although the second part of the book does depend on the first, you could jump right in at Chapter 8 if you are well-prepared and want to learn about the connection to index theory as soon as possible. You would need to refer back to earlier chapters as necessary, especially to Chapter 6. Appendix A contains various elementary results about graded algebras and modules. You should probably review it before embarking on Chapter 8.

Each chapter ends with a series of exercises. Quite a number of the exercises develop ideas not fully presented in the text; several require familiarity with other subjects (for example, geometry or representation theory); some are quite difficult. We hope that at least some of the problems are challenging enough and interesting enough to repay careful thought, since it goes without saying that working through exercises is by far the best way to learn this or any other subject.

Thanks

It is a pleasure to record our thanks to all the friends and colleagues who have helped and inspired us in the writing of this book. For guiding our mathematical paths in the direction of geometric analysis we are indebted to our doctoral advisers, Peter Fillmore and Michael Atiyah, and to the writings and encouragement of Ron Douglas. While the book was being written we benefited from discussions with many people, including Paul Baum, Martin Bridson, Alain Connes, Simon Donaldson, Pierre Julg, Gennadi Kasparov, Vincent Lafforgue, Erik Pedersen, Georges Skandalis, and Shmuel Weinberger. Dorin Dumitrascu and Nick Wright read the manuscript and made many helpful suggestions. We express our thanks, too, to all the participants in the geometric functional analysis seminar at Penn State, and to those in the topology seminar at Oxford.

We thank the Pennsylvania State University, the National Science Foundation, the University of Oxford, the University of Colorado, and the Principal, Fellows and Scholars of Jesus College, Oxford, for their support during the period of writing of this book. We also thank the staff of Oxford University Press for their patience and skill in bringing it into production.