Introduction

The two basic kinds of homology theories for algebras that are considered in *non-commutative geometry* are K-*theory* and *cyclic* (*co*)*homology*. *Periodic cyclic homology* $\mathrm{HP}_*(A)$ extends de Rham cohomology $\mathrm{H}^*_{\mathrm{dR}}(M)$ for manifolds in the sense that

$$\mathrm{HP}_n\big(\mathcal{C}^\infty(M)\big) \cong \bigoplus_{m\in n+2\mathbb{Z}} \mathrm{H}^m_{\mathrm{dR}}(M)$$

for any smooth compact manifold M; here $\mathcal{C}^\infty(M)$ denotes the Fréchet algebra of smooth functions on M. The dual theory $\mathrm{HP}^*(A)$ is called *periodic cyclic* co*homology*. Both are special cases of *bivariant periodic cyclic homology* $\mathrm{HP}_*(A, B)$. We are going to study two closely related (in fact, almost identical) variants of periodic cyclic homology called *analytic cyclic homology* HA_* and *local cyclic homology* HL_*. Both have dual cohomology theories and extend to bivariant theories.

Several of the tools we develop to study these theories are useful for other purposes as well. Therefore, even if you are not at all interested in cyclic homology, you may find that Chapters 1 and 3 and parts of Chapter 2 contain interesting ideas in functional analysis; the construction in Chapter 6 can also be applied in the context of bivariant K-theory. The Appendix contains a brief survey of some preliminaries on homological algebra in symmetric monoidal categories and constructions with differential forms.

Each chapter has its own introduction which summarises its contents. In this introduction to the whole book, we first discuss the shortcomings of periodic cyclic homology that justify the existence of analytic and local cyclic homology, and we discuss their relationship to the entire cyclic cohomology of Alain Connes. Then we discuss some of the preliminaries from functional analysis and algebra that we need to understand analytic and local cyclic homology and mention a few particularly important ideas, including excision and invariance for isoradial subalgebras.

What is wrong with periodic cyclic homology?

We get $\mathrm{HP}_*(A)$ by a limiting process from the *cyclic homology* $\mathrm{HC}_*(A)$, which is in turn computable from *Hochschild homology* $\mathrm{HH}_*(A)$ by a spectral sequence. Hence periodic cyclic homology only gives reasonable results for algebras whose Hochschild homology is sufficiently rich. Since Hochschild homology is closely related to differential forms, this requires a certain amount of differentiability. Therefore, periodic cyclic homology yields poor results for C^*-algebras; similarly, de Rham cohomology only makes sense for smooth manifolds. When we study a C^*-algebra A, we use $\mathrm{HP}_*(A^\infty)$ for an appropriate dense subalgebra A^∞, which plays the role of the subalgebra $\mathcal{C}^\infty(M)$ of smooth functions in the algebra $\mathcal{C}(M)$ of continuous functions.

This works well enough in many concrete cases. But $\mathrm{HP}_*(A^\infty)$ may depend on the choice of the subalgebra A^∞. For instance, consider a connected Lie group G, say, the circle group $\mathbb{T}^1$. Let $C^*(G)$ be the group C^*-algebra of G. Let $\mathcal{C}^\infty(G) \subseteq C^*(G)$ be the dense subalgebra of smooth functions. Let $\mathcal{R}(G) \subseteq \mathcal{C}^\infty(G)$ be the dense subalgebra that is spanned by the matrix coefficients of irreducible representations. Both $\mathcal{R}(G)$ and $\mathcal{C}^\infty(G)$ behave like 0-dimensional spaces, so that their periodic cyclic homology agrees with the 0th Hochschild homology, HH_0. The latter turns out to be different for $\mathcal{R}(G)$ and $\mathcal{C}^\infty(G)$ (see §2.4). Although $\mathcal{C}^\infty(G)$ is a more obvious choice of smooth subalgebra, $\mathcal{R}(G)$ gives better results in the sense that the Chern–Connes character

$$\mathrm{K}_*(A) \otimes_{\mathbb{Z}} \mathbb{C} \cong \mathrm{K}_*(A^\infty) \otimes_{\mathbb{Z}} \mathbb{C} \to \mathrm{HP}_*(A^\infty)$$

is an isomorphism for $A^\infty = \mathcal{R}(G)$, but *not* for $A^\infty = \mathcal{C}^\infty(G)$.

Another drawback of periodic cyclic cohomology is its built-in finite dimensionality: it only admits Chern–Connes characters for finitely summable Fredholm modules.

Entire to analytic cyclic homology and cohomology

Entire cyclic cohomology extends periodic cyclic cohomology by, roughly speaking, allowing infinite-dimensional cochains with sufficiently slow growth. The main application is to construct Chern–Connes characters for ϑ-summable Fredholm modules (the summability condition restricts the growth of the singular values of certain compact operators). The *JLO cocycle* defined in [53] (and named after its creators Jaffe, Lesniewski, and Osterwalder) provides an explicit formula for such a character. The JLO cocycle and related formulas are useful for index computations such as the proof of the local index formula of Alain Connes and Henri Moscovici ([10]); thus entire cyclic cohomology is interesting even for algebras like $\mathcal{C}^\infty(M)$ for which we know that it coincides with periodic cyclic cohomology.

Entire cyclic cohomology goes beyond Hochschild cohomology, but it does not yet address the first shortcoming of periodic cyclic cohomology mentioned above: a result of Masoud Khalkhali ([62]) yields $\mathrm{HE}^*(A) = \mathrm{HP}^*(A)$ for nuclear C^*-algebras, so that we seem to gain nothing. Nevertheless, some apparently small technical changes yield a theory with considerably better properties.

First we pass from entire cyclic *co*homology to the dual homology theory HE_* (see also [34]). To define this theory properly, we first need a category of algebras on which to define it. The growth condition for entire cyclic cochains only depends on the collection of (von Neumann) *bounded* subsets of A, which we call its *von Neumann bornology*. A *bornology* is a suitable family of subsets, called *bounded subsets*, and a *bornological algebra* is an algebra over $\mathbb{R}$ or $\mathbb{C}$ with a bornology such that the multiplication is bounded. This category of algebras is the natural domain for entire cyclic homology and cohomology. The chain complex that defines entire cyclic homology carries a bornology by construction, but it carries no obvious topology in general.

We may choose another bornology on a topological algebra than the von Neumann bornology. In fact, our *default* choice is the bornology of *precompact* subsets because this yields considerably better results. To avoid confusion with the existing entire theories, we call our theories *analytic* cyclic homology and cohomology and denote them by HA_* and HA^*. Both are defined for complete bornological algebras. When we apply them to topological algebras, it is understood that we equip them with the precompact bornology; in contrast, entire cyclic homology and cohomology use the von Neumann bornology.

Another reason to change notation is that the apparent importance of *entire functions* for entire cyclic cohomology is an artefact created by looking only at a *co*homology theory. The relevant function algebra is the algebra $\mathbb{C}(\!(t)\!)$ of *analytic power series*, that is, power series with a non-zero radius of convergence. The space of entire functions only appears because it is the *dual* space of $\mathbb{C}(\!(t)\!)$. Analytic power series in several non-commuting variables are a crucial ingredient in our conceptual approach to analytic cyclic homology.

Unlike the periodic and entire cyclic theories, analytic cyclic homology yields good results for C^*-algebras (with the precompact bornology). It is split-exact, C^*-stable, invariant under continuous homotopies, and additive for C^*-completions of infinite direct sums. Standard results related to the Universal Coefficient Theorem for Kasparov theory yield a Chern–Connes character $\mathrm{K}_*(A) \otimes \mathbb{C} \to \mathrm{HA}_*(A)$ and show that it is an isomorphism if A belongs to the bootstrap category (see §7.2). In particular, for locally compact topological spaces X we have

$$\mathrm{HA}_*\big(\mathcal{C}_0(X, \mathbb{C})\big) \cong \mathrm{K}^*(X) \otimes \mathbb{C} \cong \mathrm{H}^*_{\mathrm{c}}(X; \mathbb{C}).$$

We also construct a Chern–Connes character from the Kasparov K-homology $\mathrm{K}^*(A)$ to analytic cyclic cohomology $\mathrm{HA}^*(A)$; notice that we require no summability condition.

This should suffice to make you curious about analytic cyclic homology. But a lot of preparation is needed until we can prove these results. The good news is that many of these preliminary results are useful for other purposes as well; therefore, large parts of this book may be useful to readers with no interest in cyclic homology. The main difficulty in learning about analytic cyclic homology is that it requires a good deal of both functional analysis and homological algebra; even more, the results that we need from these areas are omitted in most textbooks for beginners.

Background from functional analysis: bornologies

Bornological vector spaces have not received much attention by functional analysts, although there are several situations where they work better than topological vector spaces. As a general rule, this happens whenever we combine (homological) algebra and functional analysis. I have used bornological vector spaces for problems in representation theory and homological algebra in [68]–[72]. We spend some time reviewing the basic theory of bornological vector spaces in Chapter 1.

Besides basic issues of functional analysis, we particularly emphasise categorical constructions. Much of this is needed in order to discuss the relationship between complete bornological vector spaces and inductive systems of Banach spaces. This relationship is crucial for the more exciting results about analytic cyclic homology mentioned above. Since inductive systems are rather abstract objects, the only elegant way to deal with them is via general methods from category theory. We use some of these techniques already in the context of bornological vector spaces – where they are not yet so crucial – to exhibit the similarities between both setups.

It is not surprising that tensor products are a crucial ingredient for our cyclic theories and therefore duely studied in Chapter 1. In addition, approximation properties play an important technical role in several proofs. In connection with tensor products, we emphasise *symmetric monoidal categories*. It has already been noticed by Guillermo Cortiñas and Christian Valqui that this is the right context for studying cyclic homology theories. When we study Hochschild homology or periodic cyclic homology for, say, Banach or complete locally convex topological algebras, then we use the completed projective tensor product and restrict attention to continuous maps. Here we study algebras in the categories $\mathfrak{CBorn}$ of complete bornological vector spaces and $\overrightarrow{\mathfrak{Ban}}$ of inductive systems of Banach spaces. In each case, the definitions are essentially the same but employ different tensor products. Careful readers should wonder which properties of the cyclic theories extend to these more exotic kinds of algebras. The notion of a symmetric monoidal category formalises the basic associativity, commutativity, and monoidal properties of tensor products that we need to define algebras and modules and do homological algebra with them. All the basic features of Hochschild, cyclic, and periodic cyclic homology extend to algebras in any $\mathbb{Q}$-linear symmetric monoidal category.

Another crucial concept from functional analysis is the notion of (*joint*) *spectral radius* for bounded subsets in bornological algebras. The spectral radius $\varrho(S)$ of a bounded subset S in a bornological algebra is defined exactly as for a single element: it is the infimum of the set of scalars $r > 0$ for which

$$(r^{-1} S)^{\infty} := \bigcup_{n=1}^{\infty} (r^{-1} S)^n$$

is bounded. This notion contains more information than the spectral radius of single elements because the elements of S need not commute. Roughly speaking, spectral radius estimates ensure that certain power series in several non-commuting variables converge.

The joint spectral radius is an intrinsically bornological concept because it deals with subsets. Even for commutative algebras, where the joint spectral radius usually contains no more information than the spectral radius for single elements, the bornological framework is more suitable to study the functional calculus. This was noticed already by Lucien Waelbroeck ([107]).

We use the spectral radius to define two classes of bornological algebras with good functional calculus. A bornological algebra is called *locally multiplicative* if

$\varrho(S) < \infty$ for all bounded subsets S, and *analytically nilpotent* if $\varrho(S) = 0$ for all bounded subsets S. Both classes of algebras play an important role for analytic cyclic homology.

The spectral radius generates a useful notion of smooth subalgebra. Let A and B be locally multiplicative bornological algebras. A bounded homomorphism $f: A \to B$ with dense range (in a suitable sense) is called *isoradial* if $\varrho\big(f(S); B\big) = \varrho(S; A)$ for all bounded subsets $S \subseteq A$. In most applications, f is injective, so that we may view A as a dense subalgebra of B. The isoradiality condition means that $A \subseteq B$ is closed under functional calculus in several non-commuting variables. We check that this condition behaves nicely with respect to various constructions like extensions and tensor products, and we study several important examples. It is much harder to get a good notion of smooth subalgebra in the context of topological algebras; the most useful definition is due to Bruce Blackadar and Joachim Cuntz (see [3]).

Some relevant algebraic notions

We study algebras and modules in the somewhat exotic categories $\mathfrak{Cborn}$ and $\overrightarrow{\mathfrak{Ban}}$. As we have already explained, we use the framework of symmetric monoidal categories to deal with this. We recall a few basic facts about algebras and modules in this generality in Chapter 1 and in the Appendix. We also need various familiar results about Hochschild homology and periodic cyclic homology; we briefly discuss them in the Appendix, mostly in the general framework of symmetric monoidal categories.

The categories $\mathfrak{Cborn}$ and $\overrightarrow{\mathfrak{Ban}}$ are not Abelian, so that we need homological algebra over non-Abelian categories. A side effect of this is that the passage from chain complexes to homology forgets too much information. For chain complexes of vector spaces, the homology functor is an equivalence of categories from the homotopy category of chain complexes to the category of vector spaces; thus we lose no information at all. For chain complexes in more general Abelian categories, things already get more complicated; but at least we can detect exactness of chain complexes using homology. Once we are in non-Abelian categories, even this fails.

The lesson is that we should not take homology and instead consider functors with values in suitable homotopy categories or derived categories of chain complexes. This is crucial in [69], [71], [72] where we need certain chain complexes to be *contractible* and not just exact. Similarly, we usually treat analytic cyclic homology as a functor HA to the category of chain complexes of bornological vector spaces. This point of view is already implicit in the definition of *bivariant* cyclic homology theories because the elements of the bivariant group $\mathrm{HA}_0(A, B)$ are exactly the morphisms $\mathsf{HA}(A) \to \mathsf{HA}(B)$ in a suitable homotopy category of chain complexes.

The most sophisticated homological algebra that we need enters in the definition of bivariant *local* cyclic homology; our definition is essentially equivalent to the original one by Michael Puschnigg in [86], [88]; but we use a more efficient category of algebras. This both simplifies and generalises the theory.

The dual theory HA^* is not as well-behaved as we would like; for instance, we cannot compute $\mathrm{HA}^*(A)$ for $A := \mathcal{C}_0\big((0,1]\big)$, although we know that $\mathsf{HA}(A)$ is *locally* contractible, that is, on each bounded subset of $\mathsf{HA}(A)$ we can define a contracting homotopy. But these maps do not fit together to a global map, and the failure of the Hahn–Banach Theorem for bornological vector spaces prevents us from computing the *co*homology. To repair this defect, we replace ordinary cohomology by an appropriate derived functor (called *local* cohomology). This yields local cyclic cohomology and bivariant local cyclic cohomology. These constructions become more transparent in the context of inductive systems of chain complexes. The basic idea is that we replace an inductive system of chain complex by its *homotopy direct limit*.

There is a variant $\mathrm{HL}_*(A)$ of $\mathrm{HA}_*(A)$ as well, which repairs the lack of exactness of the completion functor for bornological vector spaces. But this makes no difference in practice: we show in Chapter 2 that $\mathrm{HL}_*(A) \cong \mathrm{HA}_*(A)$ if A is a Fréchet algebra (with the precompact bornology) or if A has a suitable approximation property. Even though the two theories agree for all practical purposes, we must first distinguish them to prove such a statement. The assertions about the analytic cyclic homology of C^*-algebras are based on this isomorphism and therefore implicitly use the local cyclic theory.

The Cuntz–Quillen approach to cyclic theories

The chain complexes $\mathsf{HA}(A)$ and $\mathsf{HL}(A)$ are defined in Chapter 2 as completions of the cyclic bicomplex because this approach to cyclic homology is probably known to most readers. But later we switch to another chain homotopy equivalent complex that is more adequate for the analytic and local cyclic theories. The cyclic bicomplex relates periodic cyclic homology to Hochschild homology and thus provides a useful scheme to compute *periodic* cyclic homology and cohomology. But since the whole point of the analytic and local cyclic theories is to go beyond the limitations of Hochschild homology, the cyclic bicomplex offers no clues how to understand these theories.

The approach of Joachim Cuntz and Daniel Quillen ([25], [26]) is more useful because it allows to treat periodic, analytic, and local cyclic homology in a very similar fashion. We use it both to prove properties shared by all three theories and to establish the special features of the local cyclic theory.

The Cuntz–Quillen approach uses two ingredients: a completed tensor algebra and the $\mathcal{X}$-complex. The $\mathcal{X}$-complex is a very small quotient of the cyclic bicomplex, which is therefore easy to study by hand; we recall its definition in §A.6.3. The most important ingredient is the completed tensor algebra. The usual tensor algebra is defined by a universal property: algebra homomorphisms $\mathrm{T}A \to B$ correspond to linear maps $A \to B$; it is not so interesting because it completely forgets the algebra structure. The completed tensor algebras that we study have a similar universal property for linear maps that are almost multiplicative in a suitable sense.

To make this precise, we need the *curvature* of a linear map $f: A \to B$, which is the bilinear map

$$\omega_f: A \times A \to B, \quad (a_1, a_2) \mapsto f(a_1 \cdot a_2) - f(a_1) \cdot f(a_2).$$

The map f is an algebra homomorphism if and only if $\omega_f = 0$. We say that f has *analytically nilpotent curvature* if the spectral radius of $\omega_f(S, S) \subseteq B$ vanishes for each bounded subset $S \subseteq A$. We briefly call such maps *lanilcurs*. The *analytic tensor algebra* $\mathcal{T}A$ is defined so that algebra homomorphisms $\mathcal{T}A \to B$ correspond to lanilcurs $A \to B$. This universal property is our main tool. We will see that $\mathsf{HA}(A)$ is chain homotopy equivalent to $\mathcal{X}(\mathcal{T}A)$. The pro-tensor algebra that appears in the Cuntz–Quillen approach to periodic cyclic homology can be characterised by a similar universal property, which we state in Chapter 4. Various formal properties like homotopy invariance for smooth homotopies, invariance under nilpotent extensions, stability for algebras of nuclear operators, and additivity can be proved easily in this framework. Furthermore, it allows us to compute some simple examples.

The deepest common property of the periodic, analytic, and local cyclic theories is *excision*. One way to formulate this is that an extension of algebras $I \rightarrowtail E \twoheadrightarrow Q$ with a bounded linear section induces an exact triangle

$$\mathsf{HA}(I) \to \mathsf{HA}(E) \to \mathsf{HA}(Q) \to \mathsf{HA}(I)[1]$$

in the homotopy category of chain complexes; this induces various long exact sequences for homology.

Excision is crucial for many computations. It was established by Joachim Cuntz and Daniel Quillen ([27]) for the periodic cyclic theory and by the author and Michael Puschnigg for the analytic and local cyclic theories. The elegant proof we present here appeared previously in [66].

Invariance for smooth subalgebras

A remarkable property of local cyclic homology is that the obvious map

$$i: \mathcal{C}^\infty(M) \to \mathcal{C}(M)$$

for a smooth compact manifold M is an HL-equivalence. More generally, the same statement holds for any isoradial homomorphism that has approximate bounded linear sections in a suitable sense. We want to indicate why this is true. Along the way, we see why it is important to use precompact bornologies and invert local chain homotopy equivalences. The following sketch is made more precise in Chapter 6.

Any smoothing operator yields a bounded linear map $s: \mathcal{C}(M) \to \mathcal{C}^\infty(M)$. There is a sequence (s_n) of smoothing operators such that $i \circ s_n$ and $s_n \circ i$ converge towards the identity maps on $\mathcal{C}(M)$ and $\mathcal{C}^\infty(M)$; this convergence is not uniform in norm, but it is *uniform on precompact subsets*. This is the point where it is crucial to use the precompact bornology on $\mathcal{C}(M)$. Since $i \circ s_n$ converges towards the identity map, the curvature $\omega_{i \circ s_n} = i \circ \omega_{s_n}$ converges towards $\omega_{\mathrm{id}} = 0$. Thus $\varrho\big(i \circ \omega_{s_n}(K, K)\big) \to 0$; now we use that i is isoradial to conclude that $\varrho\big(\omega_{s_n}(K, K)\big) \to 0$. That is, the maps (s_n) *approximately* have analytically nilpotent curvature for $n \to \infty$.

The same argument that shows that lanilcurs induce maps on $\mathcal{T}A$ also shows that maps with sufficiently small curvature such as s_n induce maps on bounded subsets of $\mathcal{T}A$. More precisely, for each bounded submultiplicative disk $S \subseteq \mathcal{T}A$ there is $n \in \mathbb{N}$ such that s_n induces an algebra homomorphism on the Banach subalgebra of $\mathcal{T}A$ generated by S.

The composite maps $i \circ s_n$ and $s_n \circ i$ are homotopic to the identity map via an affine homotopy. Again, these homotopies have approximately analytically nilpotent curvature. As a result, the map i induces a *local* homotopy equivalence $\mathcal{T}\mathcal{C}^\infty(M) \to \mathcal{T}\mathcal{C}(M)$. Since the $\mathcal{X}$-complex for quasi-free algebras is homotopy invariant (for smooth homotopies, say), the induced chain map $\mathcal{X}\big(\mathcal{T}\mathcal{C}^\infty(M)\big) \to \mathcal{X}\big(\mathcal{T}\mathcal{C}(M)\big)$ is a local chain homotopy equivalence. The definition of the bivariant local cyclic theory ensures that such chain maps become invertible. This finishes the proof that i is an HL-equivalence.

What is missing?

Finally, there are a few important aspects of the infinite cyclic cohomology theories that we study that are left out in this book due to limitations of time, space, and energy.

First, there is the JLO cocycle ([53]) and its bivariant generalisation by Denis Perrot ([79], [80]). Instead, I present another character construction that does not require any summability condition. It would be worthwhile to compare these constructions.

Secondly, there is the work of Michael Puschnigg on the local cyclic homology of group algebras ([85], [87]). It is difficult to compute local cyclic homology by hand. Besides some easy cases, all the computations we shall do in this book use formal properties of the theory to reduce the problem to a K-theory computation. It is an important point that there are examples where local cyclic homology can be computed directly and is more accessible than K-theory.

Thirdly, Christian Voigt has extended periodic and analytic cyclic cohomology theories to the equivariant setting for algebras with actions of groups or even quantum groups in [103]–[105]. This theory is also based on suitable completed tensor algebras, so that many of our arguments carry over to it.